

Managed Switches

Ethernet Managed Switching Hubs

CTRLink®

Software Manual for Console Access

Version 5.x

Covering Product Series: EICP_M, EIDX_M, EISX_M

TD020850-0MG



CONTEMPORARY **CONTROLS**®

Trademarks

Contemporary Controls and CTRLink are registered trademarks of Contemporary Control Systems, Inc. Other product names may be trademarks or registered trademarks of their respective companies.

Copyright

© Copyright 2013, by Contemporary Control Systems, Inc. All rights reserved. No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual, or otherwise, without the prior written permission of:

Contemporary Control Systems, Inc.	Tel:	+1-630-963-7070
2431 Curtiss Street	Fax:	+1-630-963-0109
Downers Grove, Illinois 60515 USA	E-mail:	info@ccontrols.com
	WWW:	http://www.ccontrols.com

Contemporary Controls Ltd	Tel:	+44 (0)24 7641 3786
14 Bow Court	Fax:	+44 (0)24 7641 3923
Fletchworth Gate	E-mail:	ccl.info@ccontrols.com
Coventry CV5 6SP UK	WWW:	http://www.ccontrols.co.uk

Contemporary Controls GmbH	Tel:	+49 341 520359 0
Fuggerstraße 1 B	Fax:	+49 341 520359 16
04158 Leipzig, Germany	E-mail:	ccg.info@ccontrols.com
	WWW:	http://www.ccontrols.de

Disclaimer

Contemporary Control Systems, Inc. reserves the right to make changes in the specifications of the product described within this manual at any time without notice and without obligation of Contemporary Control Systems, Inc. to notify any person of such revision or change.

**WARNING — This is a Class A product as defined in EN55022.
In a domestic environment this product may cause radio interference
in which case the user may be required to take adequate measures.**

1 Table of Contents

1	TABLE OF CONTENTS.....	3
2	HISTORY	5
3	INTRODUCTION.....	5
3.1	Software.....	6
3.2	Product Series	6
4	ADVANCED OPERATION.....	7
4.1	General Considerations	7
4.1.1	LEDs.....	7
4.1.2	Initial Access	7
4.1.2.1	HyperTerminal.....	7
4.1.3	Username and Password	8
4.2	Main Menu	9
4.2.1	System Configuration	10
4.2.1.1	Configure IP Address	11
4.2.1.2	Configure Ports	12
4.2.1.3	Configure Port Trunking (Copper Ports Only)	13
4.2.1.4	Configure Port Mirroring.....	15
4.2.1.5	Configure 802.1Q VLAN.....	17
4.2.1.6	Configure Filtering and Forwarding Table	24
4.2.1.7	Configure QoS (Quality of Service).....	27
4.2.1.8	Configure Relay	33
4.2.1.9	Configure Redundancy.....	37
4.2.1.10	Configure Rate Control.....	45
4.2.1.11	Configure Port Security	48
4.2.1.12	Configure IGMP Snooping	49
4.2.2	SNMP Configuration.....	50
4.2.2.1	Configure System Information.....	51
4.2.2.2	Configure SNMP Community	52
4.2.2.3	Configure SNMP Trap Receivers	53
4.2.3	Performance Monitoring	54
4.2.3.1	Monitor Port Traffic.....	55
4.2.3.2	Browse Address Table	58
4.2.3.3	Monitor Switch History.....	59
4.2.3.4	Monitor Switch Temperature	60
4.2.3.5	Monitor STP Port Status.....	61
4.2.4	Upload/Download Settings	63

5	APPENDIX.....	64
5.1	SNMP	64
5.1.1	Managed Objects for TCP/IP Based Internet (MIB-II) — From RFC 1213 ..	64
5.1.1.1	'System' group 1.3.6.1.2.1.1.....	64
5.1.1.2	'Interfaces' group 1.3.6.1.2.1.2.....	65
5.1.1.3	'IP' group 1.3.6.1.2.1.4	67
5.1.1.4	'ICMP' group 1.3.6.1.2.1.5.....	68
5.1.1.5	'TCP' group 1.3.6.1.2.1.6	68
5.1.1.6	'UDP' group 1.3.6.1.2.1.7	69
5.1.1.7	'Transmission' group 1.3.6.1.2.1.10	69
5.1.1.8	'SNMP' group 1.3.6.1.2.1.11	69
5.1.2	Managed Objects for Bridges — From RFC 1493.....	72
5.1.2.1	'dot1dBase' group 1.3.6.1.2.1.17.1	72
5.1.2.2	'dot1dTp' group 1.3.6.1.2.1.17.4	72
5.1.2.3	'dot1dTpFdbTable' 1.3.6.1.2.1.17.4.3	73
5.1.2.4	'dot1dTpPortTable' 1.3.6.1.2.1.17.4.4	73
5.1.3	Managed Objects for Ethernet-like Interface Types — From RFC 1643	74
5.1.3.1	Ethernet-like Statistics Group — 'dot3StatsTable' 1.3.6.1.2.1.10.7.2... ..	74
5.1.4	Evolution of the Interface Group of MIB-II — From RFC 1573	76
5.1.5	Private Managed Objects	77
5.1.5.1	Relay Group – 1.3.6.1.4.1.17384.1.1.2	77
5.1.5.2	RapidRing Group – 1.3.6.1.4.1.17384.1.1.3.....	77
5.1.6	Message Format for SNMP Operations.....	78
5.1.6.1	Format of Command Messages	78
5.1.6.2	Traps for SNMPv1.....	79

2 History

- 3/19/2004 Initial Release
- 5/24/2004 Web Browser Support and VLAN ID range limits
- 7/20/2004 RapidRing Support
- 3/01/2005 Added IGMP Snooping, Rate Control & Port Security
- 9/01/2005 Added STP/RSTP
- 3/15/2013 Added VLAN content for 16- and 24-port switches; dropped CD-ROM references

3 Introduction

Managed switches in the CTRLink® family provide capabilities beyond those in both Plug and Play (PnP) and Configurable Switches. Besides conventional PnP features (auto-negotiation, 10/100 Mbps data rate, half- or full-duplex operation, flow control), a managed switch adds advanced features usually found only in high-end switches:

Rapid Spanning Tree Protocol (RSTP) provides a standardized network redundancy scheme with improved network recovery time over Spanning Tree Protocol (STP).

RapidRing™ provides high speed network redundancy — allowing recovery from a link loss in under 300 ms.

VLAN allows the physical network to be configured as multiple virtual local area networks — limiting broadcast/multicast domains and improving performance.

Trunking allows ports to be associated in groups — each group functioning as a high-speed backbone to another managed switch.

QoS provides message priority with one of these priority schemes — port-based, MAC-based, 802.1p, DiffServ, or TOS.

Rate Control allows variable data rates by port for bandwidth allocation.

Port Security limits port traffic to only those devices with listed MAC addresses.

IGMP Snooping allows multicasts to be limited to only relevant ports.

Port Mirroring copies traffic from one or more ports to a monitoring port.

Programmable Fault Relay provides a dry contact to a supervisory system if the switch senses a condition such as the loss or addition of a link.

Non-blocking wire-speed operation provides a maximum data rate of 148,810 packets per second for 100 Mbps Ethernet on all ports at full duplex.

Configuration is typically done through a console port connected to a Windows-based terminal emulation program such as HyperTerminal. Port parameters (data rate, duplex, flow control) can be pre-set via the console port or auto-negotiated. Each port supports the PAUSE function for full-duplex links, and uses the backpressure scheme for half-duplex segments.

Each switch is powered from a low-voltage AC or DC source — with redundant power terminals for backup considerations. Each unit includes attachments for either DIN-rail or panel mounting. The front panel features a power LED, a management status LED and bi-colour LEDs for the link status, activity, and data rate of each port.

This software is used in the EICP_M, EIDX_M, and EISX_M product series.

3.1 Sample Images of the Managed Switch Product Series



EICP8M-100T



EISX8M-100T/FC



EIDX24M-100T/FC



EIDX24MP-100T

4 Advanced Operation

4.1 General Considerations

Configuration is accomplished while the switch is connected to a computer running a terminal-emulation program such as Microsoft's HyperTerminal.

4.1.1 LEDs

To aid in troubleshooting, several LEDs have been provided.

Each **port LED** glows solid if a link exists, flashes to show activity and shows data rate by colour — green for 100 Mbps and yellow for 10 Mbps.

The **Power LED** glows solid green to indicate the presence of adequate power.

The **Status LED** on the switch front panel acts as a heartbeat and blinks every 5 seconds during normal operation. If a fault occurs, it blinks every second — except that EIDX models maintain the 5-second heartbeat and the LED turns **red** to indicate a fault.

Some EIDX models also have PoE LEDs described in the EIDX installation guide.

4.1.2 Initial Access

4.1.2.1 HyperTerminal

HyperTerminal users should connect via the serial port (Console Port) which uses standard EIA-232 protocol for configuring the switch. For proper communication :

- Set the Baud rate to 9600.

- Set the Data bits to 8.

- Set the Parity to None.

- Set the Stop bits to 1.

- Set the Flow control (handshaking) to None.

- Set Emulation to ANSI.

4.1.2.1.1 On-Screen Help

There are many configuration screens. At the bottom of each screen is a **Help Area** of either 3 or 4 lines. The top line will display messages in response to user activity. The next line will identify the action of various key strokes. The third line and (if needed) the fourth line display information about the general operation of the screen.

4.1.2.1.1.1 Common Keystrokes

Switch screens use the keyboard **spacebar** to toggle between option states, the **TAB** key to advance from field to field, the **backspace** key to erase keystrokes and the **cursor up/down arrow** keys to move up or down between fields.

4.1.3 Username and Password

With a null-modem cable connecting the switch to a suitable configuration device, a login screen *similar* to that of Figure 1 is displayed. However, your screen will identify your specific switch model number and the version of your firmware.

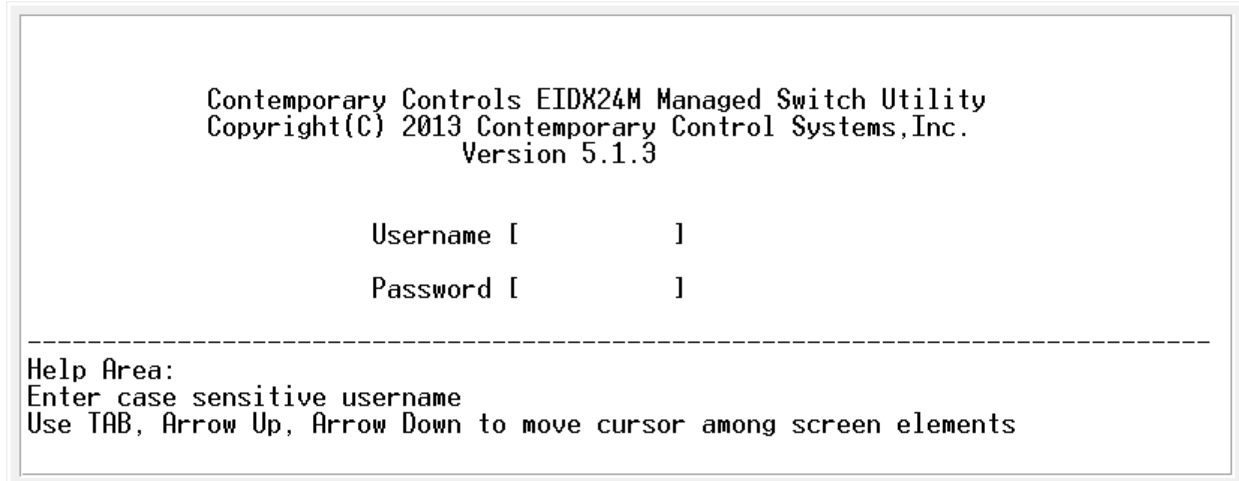


Figure 1 — Login Screen

By default, both of the *Username* and *Password* strings are “blank”. To retain a blank *Username*, press the “arrow down” button while the cursor is in the *Username* box. To use a blank *Password*, press “Enter” while the cursor is in the *Password* box.

Both *Username* and *Password* are case sensitive and the number of characters in each string can range from 0 (blank) through 10. Once these strings have been modified, future access to the management features of the switch will be denied unless the user re-enters the correct information. However, should either of these strings be forgotten, a means exists to reset them to their default values. The “reset” process consists of shorting pins 2 and 3 of the Console Port (Figure 2) for 10 seconds during boot. This will reset *ONLY* the *Username* and *Password* and *ONLY for this current login session* — the remaining parameters are not changed. After logging in, the username and password should be changed and saved to EEPROM to overwrite the forgotten username and/or password. If the username and password are not changed, the reset process must be repeated on power cycle.




Figure 2 — Restoring the Default Access Strings

After the *Username* and *Password* are accepted, the **Main Menu** (Section 4.2) appears.

4.2 Main Menu

A **Managed Switch Main Menu** *similar* to that of Figure 3 (which will identify your specific switch model and its MAC Address) offers the following options :

System Configuration	(explained in Section 4.2.1)
SNMP Configuration	(explained in Section 4.2.2)
Performance Monitoring	(explained in Section 4.2.3)
Username and Password	(explained in Section 4.1.3)
Save Settings to Non-Volatile Memory	(explained below)
Reset to Default Settings	(explained below)
Upload/Download Settings	(explained in Section 4.2.4)
Logout	(explained below)

Save Settings to Non-Volatile Memory will store any currently modified settings to the non-volatile memory contained in the switch.

NOTE: If the “Save Settings to Non-Volatile Memory” option is NOT selected after modifications have been made, any modified settings will be lost when a power cycle occurs.

Reset to Default Settings will restore all switch settings to their factory defaults. This will overwrite any previously stored settings in non-volatile memory and reboot the switch software.

Logout causes the HyperTerminal session (or that of any serial-based terminal emulation application) to restart and prompt the user for a *Username* and *Password*. After 5 minutes of inactivity, the switch will *automatically* logout. Then the login screen will appear and prompt “Press any key to log in...”. For as long as inactivity persists, this login screen will refresh every 10 seconds.

```
EISX8M Managed Switch Main Menu

System Configuration
SNMP Configuration
Performance Monitoring
Username and Password
Save Settings to Non-Volatile Memory
Reset to Default Settings
Upload/Download Settings
Logout

Switch MAC Address: 00-50-DB-00-13-3D

-----
Help Area:
Configure switch management features
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
```

Figure 3 — Main Menu

4.3 System Configuration

Each of the **System Configuration** menu selections of Figure 4 will activate additional sub-menus from among the following list :

Configure IP Address	(explained in Section 4.2.1.1)
Configure Ports	(explained in Section 4.2.1.2)
Configure Port Trunking	(explained in Section 4.2.1.3)
Configure Port Mirroring	(explained in Section 4.2.1.4)
Configure 802.1Q VLAN	(explained in Section 4.2.1.5)
Configure Filtering & Forwarding Table	(explained in Section 4.2.1.6)
Configure QoS	(explained in Section 4.2.1.7)
Configure Relay	(explained in Section 4.2.1.8)
Configure Redundancy	(explained in Section 4.2.1.9)
Configure Rate Control	(explained in Section 4.2.1.10)
Configure Port Security	(explained in Section 4.2.1.11)
Configure IGMP Snooping	(explained in Section 4.2.1.12)

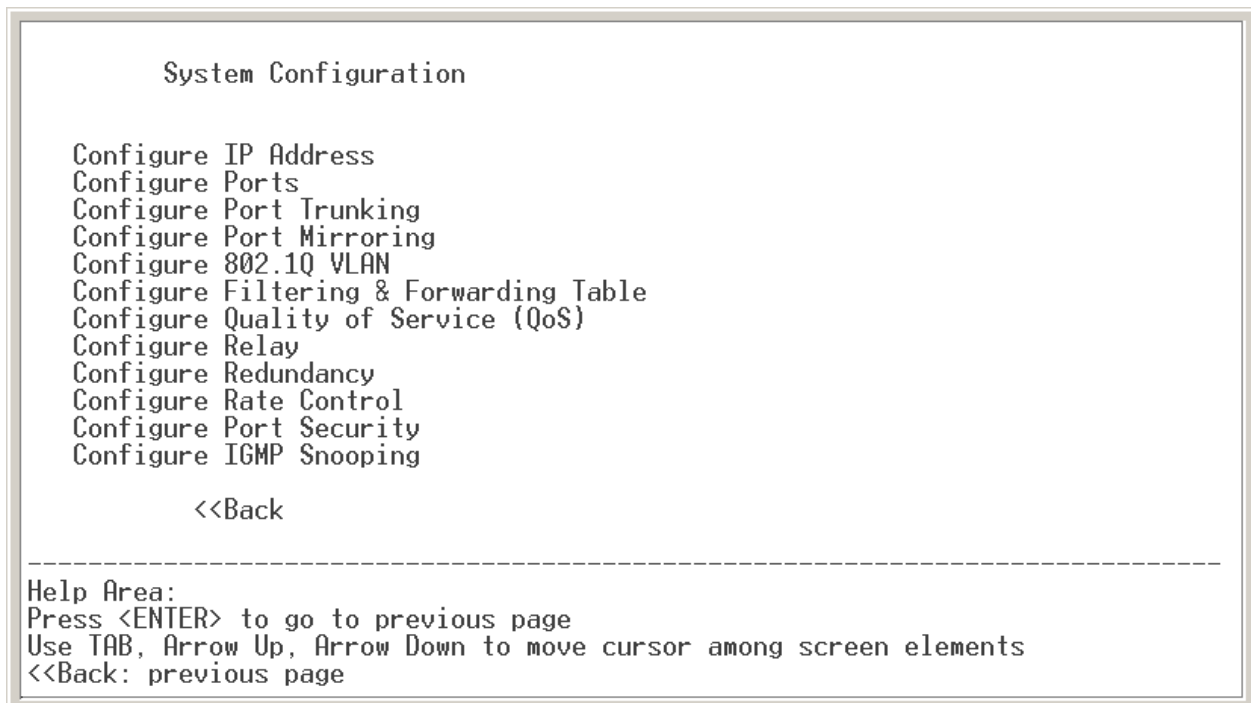


Figure 4 — System Configuration

4.3.1 Configure IP Address

Figure 5 displays the **Configure IP Address** menu with its default values. The address can be assigned in either of two ways.

By default, the switch uses a fixed IP address. And as long as the *FIXED* option is selected in the *Assigned by* field, the user must fill in the *IP Address*, *Subnet Mask* and *Default Gateway*.

The switch can have its *IP Address*, *Subnet Mask* and *Default Gateway* assigned automatically by a Dynamic Host Configuration Protocol (DHCP) server if *DHCP* is selected in the *Assigned by* field.

When the *Apply* option is exercised, the IP configuration is set to become effective and a reboot confirmation screen pops up. If the *Yes* option is exercised, the switch software restarts and the Login screen appears. If the *No* option is exercised, IP configuration aborts. If the switch is set to use a DHCP server, it will try to contact a DHCP server. If a DHCP server is not found within two minutes, the message

“No Server has been found. Set IP manually”

will display for 5 seconds and then the Login screen will appear. In this case, the user must set the *IP Address*, *Subnet Mask* and *Default Gateway* manually by using the *FIXED* option. If a DHCP server is found (usually within 20 seconds), the **Configure IP Address** screen will display the *IP Address*, *Subnet Mask* and *Default Gateway* that have been assigned by the DHCP.

```
Configure IP Address

Current IP Address:

IP Address      [192.168.92.68 ]
Subnet Mask    [255.255.255.0 ]
Default Gateway [0.0.0.0 ]

New IP Address:

Assigned by     [FIXED ]
IP Address     [192.168.92.68 ]
Subnet Mask    [255.255.255.0 ]
Default Gateway [0.0.0.0 ]

Apply    <<Back

-----
Help Area:
Assigned by DHCP or fixed IP address. Press <SPACE> to toggle
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 5 — Configure IP address Menu

4.3.2 Configure Ports

Figure 6 shows how individual Ethernet ports can be Enabled or Disabled. Also, the operating data rate or “speed” and duplex can be given a specific setting or be set to auto-negotiate.

NOTE: Certain models use fibre optic cable on two ports. For these ports, the only Speed/Duplex option is Full or Half Duplex. The Speed for fibre ports is fixed at 100 Mbps.

A single Ethernet cable will link two devices. When these devices auto-negotiate, the data rate will be 100 Mbps only if both are capable of that speed. Likewise, full-duplex will only be selected if both can support it. If only one device supports auto-negotiation, then it will match the data rate and duplex mode of the non-auto-negotiating device.

Sometimes it is advantageous to select a fixed data rate and duplex setting on both Ethernet devices to eliminate the auto-negotiation process.

When interconnecting two switches, crossover cables are normally used — but if one switch uses Auto MDIX, the communication can be via either straight-through cable or crossover cable. This functionality does not require *both* switches to have Auto MDIX.

Configure Ports					
Port	Port Status	Speed/Duplex		Auto MDIX	Flow Control
1	[Enabled]	[Auto-Negotiation]		[Enabled]	[Disabled]
2	[Enabled]	[Auto-Negotiation]		[Enabled]	[Disabled]
3	[Enabled]	[Auto-Negotiation]		[Enabled]	[Disabled]
4	[Enabled]	[Auto-Negotiation]		[Enabled]	[Disabled]
5	[Enabled]	[Auto-Negotiation]		[Enabled]	[Disabled]
6	[Enabled]	[Auto-Negotiation]		[Enabled]	[Disabled]
7	[Enabled]	[Auto-Negotiation]		[Enabled]	[Disabled]
8	[Enabled]	[Auto-Negotiation]		[Enabled]	[Disabled]

Apply <<Back Continued>>

Help Area:
Enabled or disable forwarding. Hit <SPACE> to toggle
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page

Figure 6 — Configure Ports Menu

NOTE: The sample screen above is the *first of two* displayed for a 16-port switch — or the *first of three* displayed for a 24-port switch. The succeeding screens are accessed by choosing the **Continued** option just above the dashed line. This option is not displayed in the *Configure Ports* screen of an 8-port switch.

4.3.3 Configure Port Trunking (Copper Ports Only)

Port Trunking allows two or more of ports 1–8 to be grouped with the resulting group behaving as a single logical link. The switch supports multiple trunks — each constructed of 2 or more fixed physical ports.

To keep frames in order, packets with the same source/destination MAC addressing are sent over the same trunk path — but the reverse path may follow a different link because a hash algorithm is used to balance the load between links in a trunk.

Adding more ports (links) to a trunk group will increase the communication bandwidth between two switches. Either one or two trunk groups can be defined — but always from among ports 1–8. Even with a 16- or 24-port unit, only ports 1–8 support trunking.

Port Trunking on managed switches from Contemporary Controls also provides redundancy with a fast recovery time (several milliseconds). If a link in the trunk group is lost, the remaining links immediately take over and maintain communication between the switches.

Figure 7 illustrates three two-link trunks (A, B and C) connecting two computers through four switches. This configuration could sustain a link loss in Trunk A, but within milliseconds a redundant data path would be reconstructed between the two computers. A similar recovery would manifest for a link lost in Trunk B or C. Indeed, even multiple link losses — one in each trunk group — would not disrupt communication between the two end stations except for the brief recovery time.




Figure 7 — Port Trunking

Figure 8 displays the default **Configure Port Trunking** menu. Ports to be included in a trunk would be entered with a “Y”. A trunk group must also be enabled for it to be active. If two groups are enabled, they must *not* have ports in common (overlapping ports).

Configure Port Trunking			
Group	Port-Range	Members	Status
1	1-8	[-----]	[Disabled]
2	1-8	[-----]	[Disabled]

Apply <<Back

Help Area:
Select trunk group members, (Y)=member, (-)=non-member. Hit <SPACE> to toggle
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page

Figure 8 — Configure Port Trunking

4.3.4 Configure Port Mirroring

Port mirroring allows one port to copy or “mirror” all traffic transmitted or received by one or multiple ports. This is useful when an Ethernet diagnostic tool is used. Switches generally transmit frames only to the ports involved in the conversation, therefore port mirroring is required if a diagnostic tool needs to capture network traffic. The **Configure Port Mirroring** menus allow various adjustments to how data is mirrored, but they apply *only* if the **Status** option is *Enabled*.

The **Mirror Port** is the one to which data is copied. **Ingress Mirroring Rules** apply to data *received* by the ports being mirrored. The example in Figure 9 (for an 8-port switch) has Port 1 defined as the mirror. **Source Ports** (those having their data copied) are shown by a “Y” and Figure 9 depicts Ports 2 and 3 chosen as source ports.

The **Divider** value specifies the *fraction* of messages sent by each port. For example, a Divider value of 7 will send every *seventh* message to the mirror port, while a value of 1 (as in Figure 9) will send *each* message to the mirror port.

The **MAC Address** options are :

All — mirrors *all* traffic regardless of MAC address. This is the default. Since no addressing is used, a **Source/Destination Addr** line will *not* appear.

If either of the following options are chosen, a **Source/Destination Addr** line will appear beneath the **MAC Address** line.

Source — mirrors *only* messages whose *source* addresses match the MAC Address entered by the user in the **Source Addr** below.

Destination — mirrors *only* messages whose *destination* addresses match the MAC Address entered by the user in the **Destination Addr** below.

```
Configure Port Mirroring
Status          [Disabled]
Mirror Port     [1 ] (1-8)
Ingress(Incoming) Mirroring Rules:
Source Ports    [-YY-----] (1-8)
Divider         [1  ] (1-1023)
MAC Address     [Source   ]
Source Addr     [0050DBFF0000]
                Apply    <<Back   Continued>>
-----
Help Area:
Press <ENTER> to apply new settings
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 9 — Configure Port Mirroring Ingress Rules

Egress Mirroring Rules pertain to data *transmitted* by the ports being mirrored. The example screen of Figure 10 shows that Port 4 has been designated as the only **Source Port**. Since the **Divider** value is 10, only one out of every 10 messages transmitted by Port 4 will be copied to the **Mirror Port**. The example also shows that, because the **MAC Address** setting is *All*, the source/destination address information in the copied messages will be ignored. And because the **MAC Address** in Figure 10 is set to *All*, the **Source/Destination Addr** line which appears in Figure 9 is not displayed.

```
Configure Port Mirroring (Continued)
Egress(Outgoing) Mirroring Rules:
Source Ports  [---Y----] (1-8)
Divider       [10  ] (1-1023)
MAC Address   [All   ]

Apply      <<Back

-----
Help Area:
Press <ENTER> to go to previous page
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 10 — Configure Port Mirroring Egress Rules

For 16- and 24-port switches, Figure 9 and Figure 10 would display the applicable number of spaces within the square brackets to the right of **Source Ports**.

4.3.5 Configure 802.1Q VLAN

A VLAN (Virtual Local Area Network) is comprised of devices grouped on some basis other than geographic location (i.e., by work group, security level, user type, or application). The devices *logically* behave as if tied to the same wire although they may be physically located on very different LAN segments. VLANs are configured with software, which offers much greater flexibility than hardware configuration.

A chief advantage of VLANs is that they block *broadcasts* and *multicasts* from non-VLAN ports. Most switches tend to transmit *unicast* frames sent only to ports involved in a conversation (directed messages) and cannot accommodate broadcast or multicast frames. VLANs keep broadcasts and multicasts within a VLAN group.

Another advantage of VLANs is that despite being physically relocated, a device can remain in the same VLAN — with no hardware reconfiguration needed. The VLAN supervisor can change/add workstations and manage load-balancing (bandwidth) far more easily than with a LAN modified only by hardware. Management software maintains a virtual image of how the logical and physical networks compare.




Figure 11 — VLANs

4.3.5.1 All Ports Should Be VLAN Ports

When VLANs are enabled on the switch, each port should be assigned to one or more VLANs. Such ports are called VLAN ports. If a port is not assigned to a VLAN while VLANs are enabled, that port **cannot receive messages** from the switch. A frame received from a VLAN port will only be forwarded to those ports with which it shares a VLAN membership. If the destination belongs to another VLAN, the frame will be discarded. This topology allows networks to share a common server or router, but use different VLANs for security or performance reasons.

4.3.5.2 VLAN Tags and VLAN Identifiers (VIDs)

Each VLAN frame contains an 802.1Q VLAN tag having a *VID* (VLAN Identifier) indicating to which VLAN this message belongs. The switch can be configured to allow frames with specific VID values to be received on specific ports within a VLAN.

VID values can range from 1 to 4094 — but only *within one contiguous block* of 512 values. The allowable VID blocks (ranges) are:

1–511	1024–1535	2048–2559	3072–3583
512–1023	1536–2047	2560–3071	3584–4094

Packets having VID values outside of the one defined block will be dropped.

4.3.5.3 Two Types of VLANs

The managed switch supports two types of VLANs, Port VLAN and 802.1Q VLAN. A **Port VLAN** (Section 4.2.1.5.7) is normally used to interconnect VLAN-unaware devices (such as desktop computers) which do not use VLAN tags. But **802.1Q VLANs** require 802.1Q tags in the frames passing through the switch.

4.3.5.4 Core Switches and Edge Switches

A **core switch** is connected only to devices that are VLAN aware — thus, all frames received by a core switch should *already contain* 802.1Q VLAN tags.

An **edge switch** *adds* VLAN tags to frames sent by non-VLAN aware devices and it *removes* VLAN tags from frames destined for non-VLAN aware devices.

To function in VLAN mode, the switch *requires* VLAN tags. When it performs as a *Port VLAN* switch (connected only to non-VLAN aware devices that do not use VLAN tags), a *default tag* will be applied to the untagged frames entering the switch. When these frames leave the switch, the tags should be removed. Thus, the switch can act as either a core switch or an edge switch — on a port-by-port basis. This functionality allows the switch to isolate non-VLAN aware devices by tags and give added security.

4.3.5.5 Configure 802.1Q VLAN menu

Figure 12 displays the four functions of the **Configure 802.1Q VLAN** menu :

Status (read-only) reports whether or not VLANs are enabled on the switch.
Configure Non-802.1Q Frame Drop Rules (explained in Section 4.2.1.5.9)
Configure VLAN Groups and VID (explained in Section 4.2.1.5.7)
Configure 802.1Q VLAN Tag (explained in Section 4.2.1.5.8)

```
Configure 802.1Q VLAN
Status          [Disabled]
Configure Non-802.1Q Frame Drop Rules
Configure VLAN Groups and VID
Configure 802.1Q VLAN Tag

Apply    <<Back

-----
Help Area:
Press <ENTER> to save changes. New settings will be applied after restart.
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 12 — Configure 802.1Q VLAN

4.3.5.6 VLANs and Trunking

A problem can occur if one of the ports in a trunk is not in the same VLAN as the other ports in the trunk. Messages will travel over one of the ports in a trunk, but this port cannot be selected by the user. Messages could also pass successfully through a trunk via a port in the VLAN — but if that link becomes disrupted, messages would pass through another trunk port that may *not* be in the same VLAN. Therefore *all* of the ports in a trunk must be in the *same* VLAN.

Therefore, the rule when using trunking and VLANs is for *all* ports in the trunk to be in the *same* VLAN and have the *same* default VID number — and every port in potential use should have a VLAN defined for it.

4.3.5.7 Creating VLAN Groups and VIDs

Nine VLANs can be defined using the three **Configure 802.1Q VLAN Group and VID** screens. The first screen (in this example, for an 8-port switch) appears in Figure 13 — where each group by default has a unique *VID*, contains no *Members*, is *Disabled* and has its *Tag Filter* set (Y) to remove VLAN tags from all ports as they forward messages out of the switch. To access each screen in succession, select the *Continue* option in the lower part of each of these three screens.

When creating a VLAN, several steps are required to complete the configuration. Do **not** set any group's *Status* to *Enabled* until all other configuration is complete. (This is needed to avoid illegal VID values that would result unless all VIDs have been assigned to the same continuous block of VID values as explained in Section 4.2.1.5.2).

To define the *Members* of each group, toggle the “-“ symbol into a “Y” symbol for each port in the group. To the right of the bracketed group of *Members* (up to 24 ports on some switch models) a lone “-“ symbol represents the “Management Port”. Include this port in any group for which managed port behaviour is desired.

Beneath the *Members* group is the *Tag Filter* definition for the ports in question. When a VLAN message leaves the switch, care must be exercised regarding the VLAN tag contained in the message. In Port VLAN mode, the switch will insert a VLAN tag into any message arriving from a non-VLAN-aware device. When this message exits the switch, the tag will still be present. This tag could cause a problem if the receiving device is not VLAN-aware. To remove the tag from the message, enable (set to “Y”) the *Tag Filter* entry for *Member* ports that are connected to non-VLAN aware devices.

```
Configure 802.1Q VLAN Group and VID
Group  VID  Members (1-8, M)  Status
Tag Filter (1-8)
1      [1  ]  [-----]_[-]  [Disabled]
          [YYYYYYYYY]
2      [2  ]  [-----]_[-]  [Disabled]
          [YYYYYYYYY]
3      [3  ]  [-----]_[-]  [Disabled]
          [YYYYYYYYY]

Apply    <<Back  Continue>>
-----
Help Area:
Press <ENTER> to apply new settings
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page. Note: M = Management Port.
```

Figure 13 — Configure VLAN Group and VID

4.3.5.7.1 Management Port

If a VLAN is to have full management functionality for *any* of its ports, it must include the Management Port as one of its members. This port accomplishes the management logic solely within the Ethernet controller chip and outside that chip it has no physical presence at all. It extends management functionality only to the ports with which it communicates. However, to bestow added security to its member ports, the user may wish to deliberately create a VLAN that *excludes* the Management Port — but such a group cannot support IGMP Snooping, the web server or SNMP.

4.3.5.8 Configure 802.1Q VLAN Tag

802.1Q VLANs require 802.1Q tags in the Ethernet frames that travel through the switch. With this functionality, the managed switch can participate in a VLAN as either a core switch or an edge switch.

Frames arriving from non-VLAN devices are “tagless” and cannot function in a VLAN unless they have tags added to them. The insertion of tags is a function of the ingress port through which these tagless frames pass into the switch. For this purpose, the ingress port is assigned a value of *Default Tag* that is inserted .

Figure 14 displays the **Configure 802.1Q VLAN Tag** screen. However, the default VID tag of the ingress port through which they pass must match the VID of the group to which the frames are destined. The VLAN group must also be enabled.

```
Configure 802.1Q VLAN Tag

Port      Default Tag
 1         [15 ]
 2         [15 ]
 3         [15 ]
 4         [15 ]
 5         [15 ]
 6         [15 ]
 7         [15 ]
 8         [15 ]
 M         [15 ]

Apply     <<Back

-----
Help Area:
Press <ENTER> to apply new settings
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page. Note: M = management port.
```

Figure 14 — Configure 802.1Q VLAN Tag

4.3.5.8.1.1 Example of VLAN Configuration

Example : To create a Port VLAN group with a VID of 1 and containing Ports 1–3, the VLAN must not only be defined, but Ports 1–3 in should also have their default tags set to 1. When using Port VLAN, the switch should also perform tag filtering or removal on all ports (Figure 14). In this way, applied tags will not reach VLAN unaware devices.

A special type of Port VLAN is the overlapping Port VLAN where one port is shared among multiple VLANs.

Example : Three computers on Ports 1–3 need to share a printer on Port 4, but the communication among these computers needs to be blocked. One way to accomplish this is to create four VLAN groups in the **Configure 802.1Q VLAN Group and VID** menu (Figure 13) as follows :

Group 1 : VID = 1, members 1 & 4

Group 2 : VID = 2, members 2 & 4

Group 3 : VID = 3, members 3 & 4

Group 4 : VID = 4, members 1, 2, 3 & 4

Then in the **Configure 802.1Q VLAN Tag** menu (Figure 14), the “Egress Tag Filter” would be enabled for ports 1–4 and the following VIDs would be defined :

Port 1 : 1

Port 2 : 2

Port 3 : 3

Port 4 : 4

4.3.5.9 Configure Non-802.1Q Frame Drop Rules

The managed switch supports the ability to drop non-802.1Q frames (frames without VLAN tags). As the 16-port display of Figure 15 illustrates, the switch can drop all non-802.1Q frames on a port-by-port basis. This is a useful feature for core switches.

When **Drop VID Violation Frame** is *Enabled*, each frame's VID tag will be examined to assure that the ingress port that will pass the frame belongs to the group using this tag. If it does not, the frame will be dropped. This feature can add extra security because a correct VID value does not guarantee a frame's travel through the switch. The ingress port must also belong to the defined group to pass the frame through the switch.

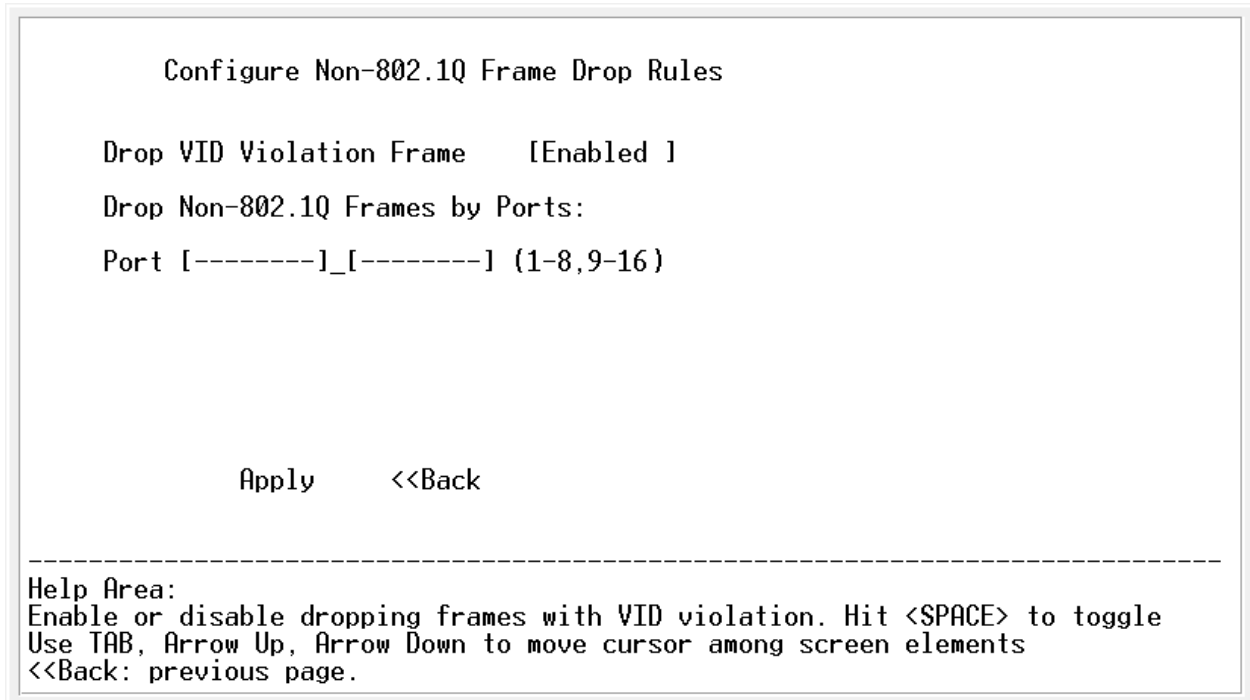


Figure 15 — Configure Non-802.1Q Frame Drop Rules

Port VLAN can be used in two different ways. In a network of devices that do not support 802.1Q, the switch can add appropriate tags to incoming messages. This will isolate the network since communication will be limited to devices in the same group or groups using the same VID. In this mode it is advisable to remove the VLAN tags on all outgoing (egress) messages (see Figure 13). However, by leaving the VLAN tags in the outgoing messages one can allow non-802.1Q devices to participate in a 802.1Q VLAN network. This second method would make the switch act as a VLAN translator for non-VLAN compliant devices .

If the network is 802.1Q compliant, one must consider whether the unit is acting as a *core switch* (in the middle of a VLAN) or as an *edge switch* (connected to non-VLAN aware devices). If the unit is acting as a core switch, the VLAN tags should not be filtered from the message. If the unit is performing as an edge switch, it should remove the VLAN tags from those ports that connect to non-VLAN aware devices. It is possible for the unit to act as *both* a core switch and as a edge switch on a port-by-port basis.

4.3.6 Configure Filtering and Forwarding Table

An Ethernet switch learns which devices are tied to which ports by monitoring the traffic carried through the switch. This information (MAC addresses and the ports to which they are attached) is stored in the **address table** which holds up to 4000 address/port associations. The switch will only transmit traffic destined for a *registered* address via the specific port associated with that address. The table behaviour follows.

During frame reception, the frame's destination address is compared with table entries for a possible match. If a match is found, the port associated with the address is noted from the table and the frame is directed to that port only. If no match is found, the frame is flooded (transmitted) to all ports.

During the reception of a *unicast* frame, the frame's source address is also compared with table entries for a match. If a match is not found, the unregistered address (and the port by which it arrived) will be added to the table. However, the address will NOT be saved if it is from the Management Port, the frame has an error or is illegal in length or if the table has no room for the new entry.

If a port-device association is not refreshed within the address table's "MAC aging time", its information will be discarded. Figure 16, displays the typical MAC aging time of 300 seconds, but this can be set as high as 1048575 seconds (over 12 *days*).

The menu displayed in Figure 16 also provides options to display additional windows for adjusting the multicast filtering and static filtering table settings :

- Configure Multicast Filtering Table** (explained in Section 4.2.1.6.1)
- Configure Static Forwarding Table** (explained in Section 4.2.1.6.2)

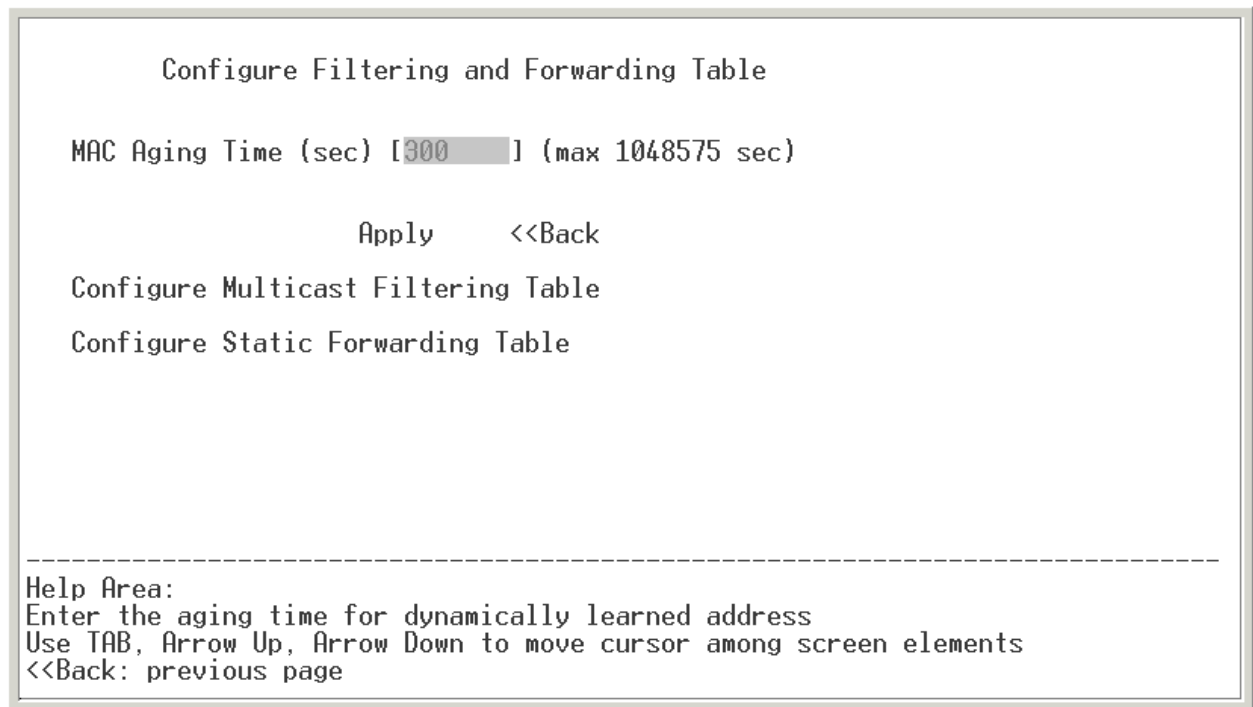


Figure 16 — Configure Filtering and Forwarding Table

4.3.6.1 Configure Multicast Filtering Table

A multicast message is one destined for two or more Ethernet devices. By default, the switch transmits such a message over all of its ports. However, the switch can filter up to ten multicast addresses so that messages sent to these addresses will only exit the switch via certain designated ports. The switch also supports IGMP Snooping which allows automatic filtering of multicast messages for devices that support multicasting as described in Section 4.2.1.12.

For an 8-port switch, these settings are illustrated in Figure 17. In this screen the user must enter the multicast **MAC Address** to represent each multicast group, the **Priority** of those messages and the **Ports** that will carry messages to that group. By default, each address is assigned a **Priority** of “Low”. The **Action** field can be toggled between *Add* and *Delete*.

For 16- and 24-port switches, more port spaces would be displayed to the right of the **Priority** listing in each **MAC Address** line.

The example of Figure 17 shows how two MAC Addresses were defined. Ports 1–5 have been associated with address 07892a310000 with a Low Priority. Ports 1 and 2 have been assigned to address 01a0c2000000 and are assigned a Medium Priority.

```
Configure Multicast Filtering Table

Action  MAC Address  Ports (1-8)  Priority
[Add   ] [01a0c2000000] [YY-----] [Medium]

Apply   <<Back

Multicast Filtering Entries (Maximum 10):  2

MAC Address  Priority  Ports (1..8) in Multicast Group
07892a310000 Low      YYYYYY---
01a0c2000000 Medium   YY-----

-----
Help Area:
Press <ENTER> to apply new settings
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page.
```

Figure 17 — Configure Multicast Filtering Table

NOTE: For an address to be accepted as a valid *multicast* address, its second digit must be *odd*. If the second digit is *even*, an error is reported in the Help Area.

4.3.6.2 Configure Static Forwarding Table

The forwarding (address) table can hold not only learned addresses, but also up to 30 *static unicast* addresses. (The sample screen of Figure 18 displays three addresses set on an 8-port switch.) Static addresses perform as if learned, but are not subject to the aging process. One of four levels of priority can also be applied to messages containing the entered MAC address of the destination.

The **Action** field can be toggled between *Add* and *Delete*. The **MAC Address** field requires a valid input from the user; it is not case sensitive. To identify the port being defined, enter a value in **Egress Port** field. The default **Priority** value is *Normal* — use the space bar to select a *Low*, *Medium* or *High* level. When the cursor is positioned in the **Apply** field, pressing the “Enter” key will complete the update.

To modify a defined address, first delete it — then redefine it with new values.

```
Configure Static Forwarding Table

Action   MAC Address   Egress Port(1-8)  Priority
[Add ]   [00ab23330001] [4 ]              [Normal]   Apply  <<Back

Static MAC Entries (Maximum 30): 3

Number  MAC Address   Egress Port(1-8)  Priority
1       0050cc450000  1                 High
2       00345b330010  1                 Normal
3       00ab23330001  4                 Normal

-----
Help Area:
Press <ENTER> to apply new settings
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 18 — Configure Static Forwarding Table

NOTE: For an address to be accepted as a valid *unicast* address, its second digit must be *even*. If the second digit is *odd*, an error is reported in the Help Area.

4.3.7 Configure QoS (Quality of Service)

In addition to the MAC-based priority applied in multicast filtering and static forwarding, the switch can assign other types of priority to its traffic to achieve various levels in what is known as Quality of Service (QoS). It can do this regardless of frame content (Port QoS) or by examining the content of every frame received by a port and assigning priority based on the port of origin. The default screen displayed in Figure 19 shows that **QoS Status** must be *enabled* before any (or all) of these methods can be applied. Below the **QoS Status** line, there are links to three configuration subscreens — one for each type of QoS. Once QoS is enabled, the individual ports can be configured by selecting the menu item :

Configure Port QoS	(explained in Section 4.2.1.7.1)
Configure 802.1p Priority	(explained in Section 4.2.1.7.2)
Configure TOS/DiffServ Priority	(explained in Section 4.2.1.7.3)

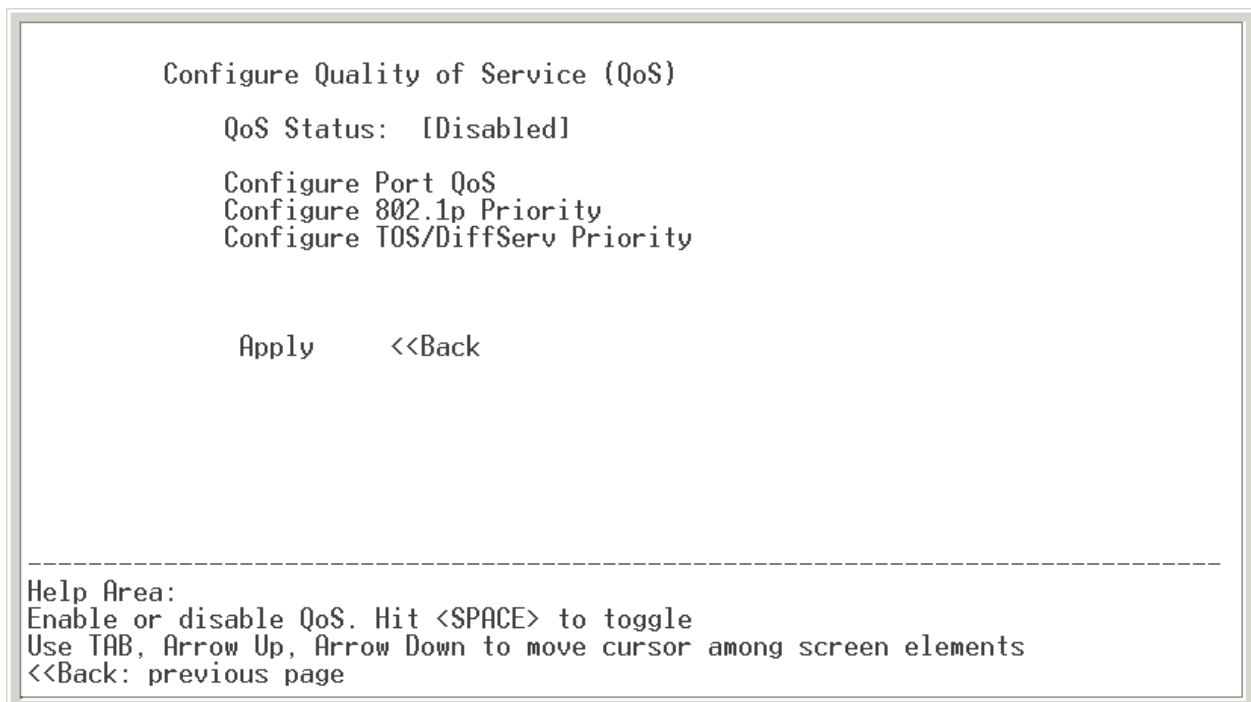


Figure 19 — Configure Quality of Service (QoS)

Although not recommended, all types of priority can be active simultaneously — giving rise to conflicts. The following hierarchy shows how these conflicts are resolved:

1. If *Port QoS* is enabled, apply its rules — otherwise,
2. If *TOS/DiffServ Priority* is enabled, apply its rules — otherwise,
3. If *802.1p Priority* is enabled, apply its rules — otherwise,
4. Apply the MAC-based priority used in multicast filtering and static forwarding.

4.3.7.1 Configure Port QoS

When QoS is enabled, **Flow Control** for each port can be enabled or disabled. For QoS to be most effective, it is recommended that each port have its flow control disabled (the default setting). When a port is operating in half-duplex mode, flow control is accomplished with backpressure. In full-duplex mode, flow control is accomplished through the PAUSE protocol. These methods and protocols can affect the ability of a port to deliver messages and can cause other messages to be delayed.

In the example of Figure 20, Port 2 has been configured for High priority while all other ports are set to Normal. But Port 1 (set to Normal priority) is the only port on which Flow Control will function because only it has had Flow Control Enabled.

```
Configure Port QoS

Port  Flow Control  Port Priority
1     [Enabled ]    [Normal]
2     [Disabled]    [High ]
3     [Disabled]    [Normal]
4     [Disabled]    [Normal]
5     [Disabled]    [Normal]
6     [Disabled]    [Normal]
7     [Disabled]    [Normal]
8     [Disabled]    [Normal]

Apply    <<Back    Continued>>

-----
Help Area:
Enabled or disable flow control. Hit <SPACE> to toggle
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 20 — Configure Port QoS

NOTE: The sample screen above is the *first of two* displayed for a 16-port switch — or the *first of three* displayed for a 24-port switch. The succeeding screens are accessed by choosing the **Continued** option just above the dashed line. This option is not displayed in the *Configure Port QoS* screen of an 8-port switch.

4.3.7.2 Configure 802.1p Priority

The IEEE 802.1p extension of IEEE 802.1Q prioritizes traffic at the data-link/MAC layer through a 3-bit header field that was never articulated in the original VLAN standard. IEEE only *suggests* 802.1p definitions; it does not mandate them.

Figure 21 shows that 802.1p priority is applied individually to ports — and just above the dashed line a link (**Map 802.1p Priority**) exists to the screen (Figure 22) where 4 priority queues allow the 8 *tags* to be mapped in various schemes. These two sample screens are for an 8-*port* switch. More ports would be displayed for a 16- or 24-*port* switch.

```
Configure 802.1p Priority
Port  802.1p
  1   [Enabled ]
  2   [Enabled ]
  3   [Enabled ]
  4   [Enabled ]
  5   [Enabled ]
  6   [Enabled ]
  7   [Enabled ]
  8   [Enabled ]

Apply  <<Back  Map 802.1p Priority>>
-----
Help Area:
Enable/Disable 802.1p priority. Hit <SPACE> to toggle
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 21 — Configure 802.1p Priority

```
Map 802.1p Priority
Tag   Priority Queue
  0   [Low   ]
  1   [Low   ]
  2   [Normal]
  3   [Normal]
  4   [Medium]
  5   [Medium]
  6   [High  ]
  7   [High  ]

Apply  <<Back
-----
Help Area:
Queue with priority Low/Normal/Medium/High. Hit <SPACE> to toggle
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 22 — Map 802.1p Priority

4.3.7.3 Configure TOS/DiffServ Priority

The IP header contains an eight-bit field originally known as the **Type of Service (TOS)** field — but TOS priority had little acceptance. Subsequently, these eight bits were redefined to become the more popular **Differentiated Services (DiffServ)** field.

When configuring the TOS/DiffServ priority (Figure 23), you must choose TOS or DiffServ. This sample screen is for an 24-port switch.

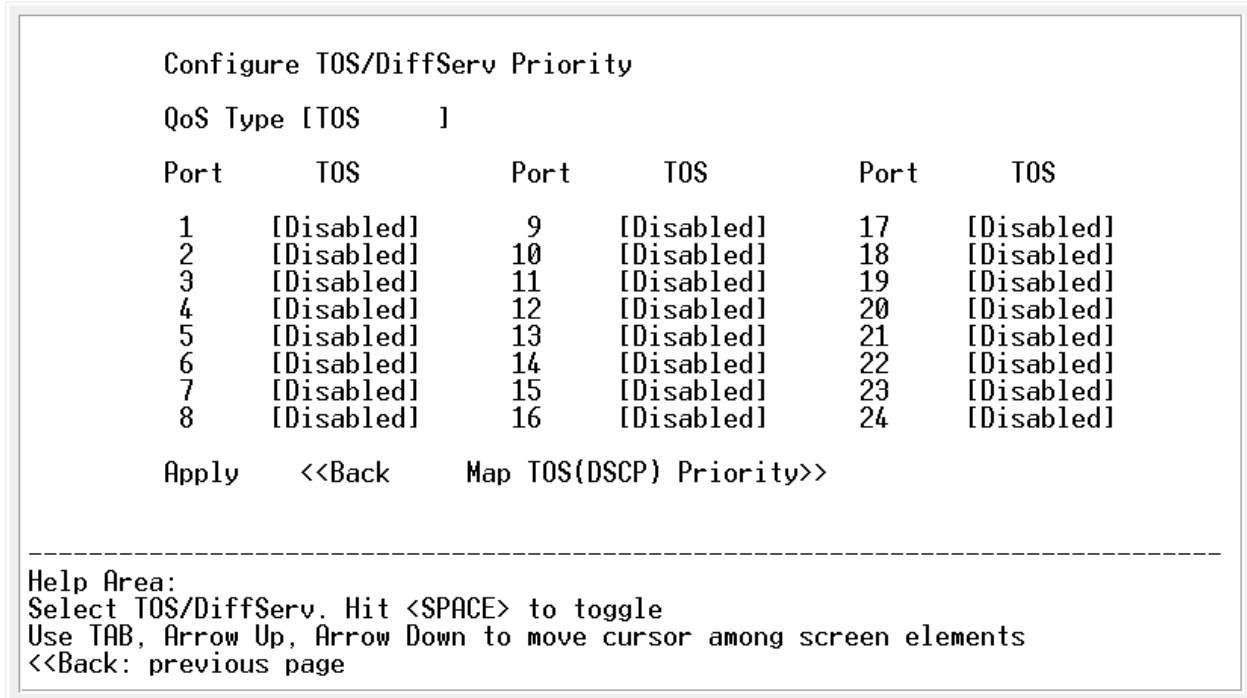


Figure 23 — Configure TOS/DiffServ Priority

You register your choice of QoS in the topmost field which is labelled **QoS Type** (the default setting is *TOS*). Below the TOS/DiffServ selection field are the fields for enabling or disabling the individual ports. Fewer ports would be displayed for an 8- or 16-port switch. Regardless of the type of QoS chosen, each port (*Disabled* by default) must be individually *Enabled* to establish its QoS service.

If TOS priority is selected, one additional screen is used for mapping the priority details, but two more screens are needed if the selected QoS Type is DiffServ. Access to these screens is via the *Map TOS(DSCP) Priority* option in the lower right portion of the screen. Each screen is discussed and illustrated below.

4.3.7.3.1 Map TOS Precedence and Priority Queue

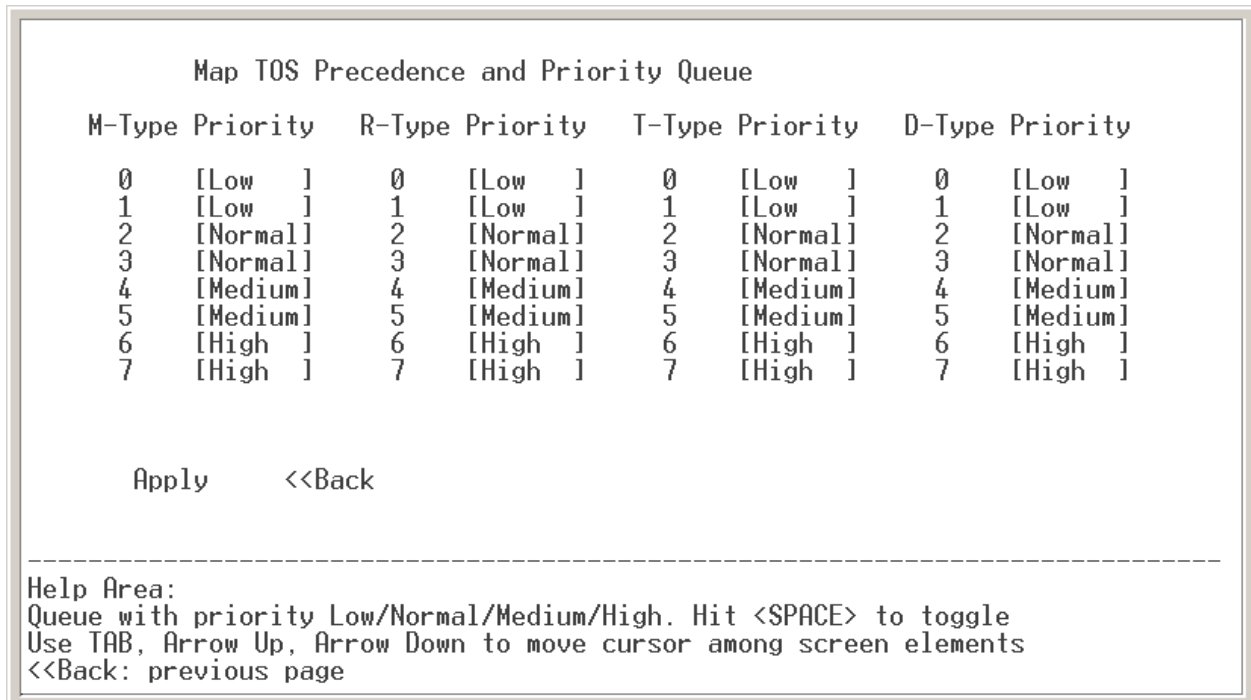


Figure 24 — Map TOS Precedence and Priority Queue

Figure 24 displays the default TOS parameters that will apply if you have chosen TOS as your QoS type. Although TOS priority is supported by few TCP/IP implementations, it is provided as a QoS option in managed switches from Contemporary Controls. One of the earliest methods of QoS for Internet Protocol, TOS uses the second octet (the TOS field) of the IP frame header — as described in RFC791 and RFC1349. The first three bits of this octet set the priority (*Precedence*). The next four bits (known as the *TOS* bits) define the tradeoffs among these four service objectives:

- minimize **m**onetary cost (*M-Type Priority*) *
- maximize **r**eliability (*R-Type Priority*)
- maximize **t**hroughput (*T-Type Priority*)
- minimize **d**elay (*D-Type Priority*)

* This is sometimes referred to as “C-Type” priority or type of service.

The final bit of the TOS field was unused until DiffServ was defined.

4.3.7.3.2 Map DiffServ DSCP and Priority Queue

Differentiated Services (DiffServ) was first described by RFC2474 which redefines the TOS octet's first six bits as the **Differentiated Services Code Point (DSCP)**. This field selects the per-hop behaviour (PHB) — how packets are queued at network nodes. RFC 2475 describes DiffServ methods for scalable differentiated services on the Internet.

Figure 25 displays the first screen of default parameters that apply if you chose DiffServ as your QoS type. Figure 26 displays the remainder of the default parameters.

```
Map DiffServ DSCP and Priority Queue

DSCP Priority   DSCP Priority   DSCP Priority   DSCP Priority
   0 [Low  ]     1 [Low  ]     2 [Low  ]     3 [Low  ]
   4 [Low  ]     5 [Low  ]     6 [Low  ]     7 [Low  ]
   8 [Low  ]     9 [Low  ]    10 [Low  ]    11 [Low  ]
  12 [Low  ]    13 [Low  ]    14 [Low  ]    15 [Low  ]
  16 [Normal]   17 [Normal]  18 [Normal]  19 [Normal]
  20 [Normal]   21 [Normal]  22 [Normal]  23 [Normal]
  24 [Normal]   25 [Normal]  26 [Normal]  27 [Normal]
  28 [Normal]   29 [Normal]  30 [Normal]  31 [Normal]

Apply    <<Back  Continued>>

-----
Help Area:
Queue with priority Low/Normal/Medium/High. Hit <SPACE> to toggle
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 25 — Map DiffServ DSCP and Priority Queue (Part 1)

```
Map DiffServ DSCP and Priority Queue

DSCP Priority   DSCP Priority   DSCP Priority   DSCP Priority
  32 [Medium]   33 [Medium]   34 [Medium]   35 [Medium]
  36 [Medium]   37 [Medium]   38 [Medium]   39 [Medium]
  40 [Medium]   41 [Medium]   42 [Medium]   43 [Medium]
  44 [Medium]   45 [Medium]   46 [Medium]   47 [Medium]
  48 [High  ]   49 [High  ]   50 [High  ]   51 [High  ]
  52 [High  ]   53 [High  ]   54 [High  ]   55 [High  ]
  56 [High  ]   57 [High  ]   58 [High  ]   59 [High  ]
  60 [High  ]   61 [High  ]   62 [High  ]   63 [High  ]

Apply    <<Back  Continued>>

-----
Help Area:
Press <ENTER> to apply new settings
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 26 — Map DiffServ DSCP and Priority Queue (Part 2)

4.3.8 Configure Relay

The switch has a relay output that can be used to signal the occurrence of one or more events. The relay is used to indicate the loss of a link or presence of a link on one or many ports. The **Configure Relay** screen (Figure 27) has three menu items for more screens.

Configure Relay Features	(explained in Section 4.2.1.8.1)
Configure Port Link Monitoring	(explained in Section 4.2.1.8.2)
Display Faults	(explained in Section 4.2.1.8.3)

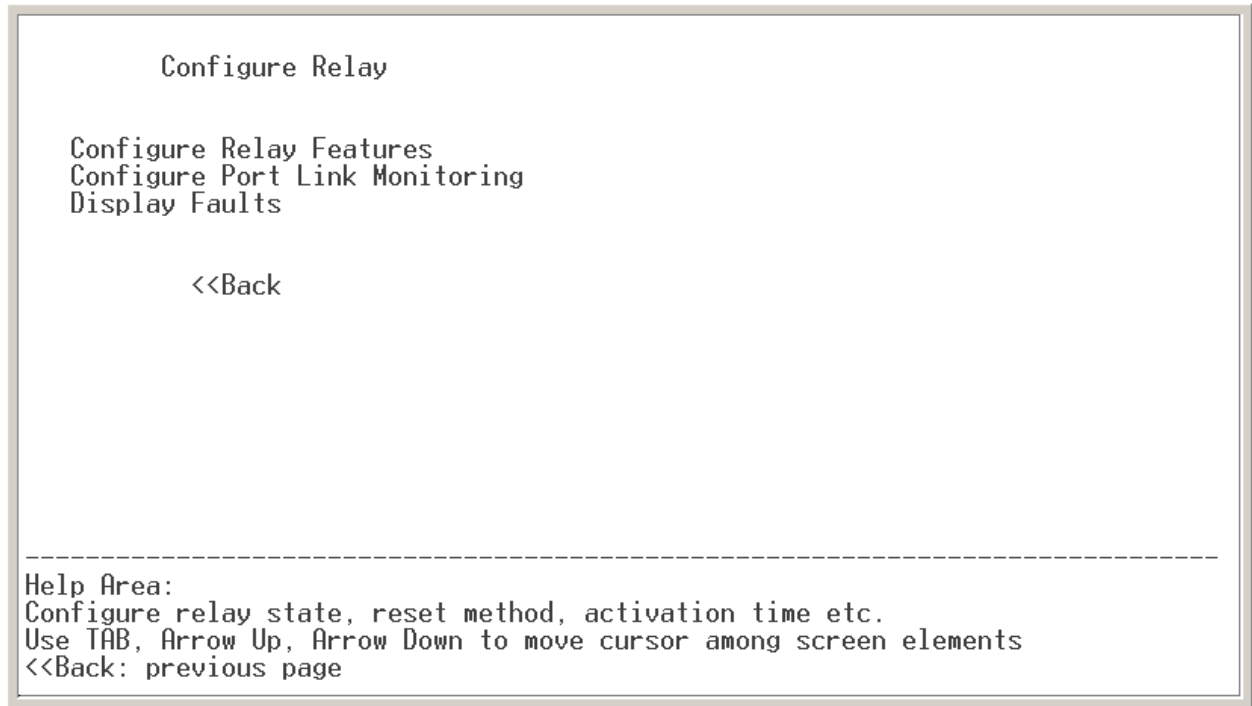


Figure 27 — Configure Relay

4.3.8.1 Configure Relay Features

Figure 28 displays the following four features :

Port Monitoring (link monitoring of specific ports) can be either enabled or disabled.

Relay State determines the behaviour of the relay. By selecting “Make on Fault”, it will *close* its contacts once a fault is detected. “Break on Fault” will cause the relay to normally keep its contacts closed and *open* them upon detection of a fault. . A fault is active when the condition of a monitored link or port matches its monitored state (see 4.2.1.8.2).

Relay Activation After Power specifies a port-monitoring delay which allows the switch to stabilize for 1 to 999 seconds after power-up. This is provided because, after power-up several seconds may be required for the switch to complete auto-negotiation of the data rate and duplex mode for each port. If the relay were not inhibited during this time, it could repeatedly activate without a true fault existing.

Relay Reset Method is automatic by default, but can be set to manual. Selecting “Manual” will cause a new menu item to appear that will need to be chosen to reset the relay.

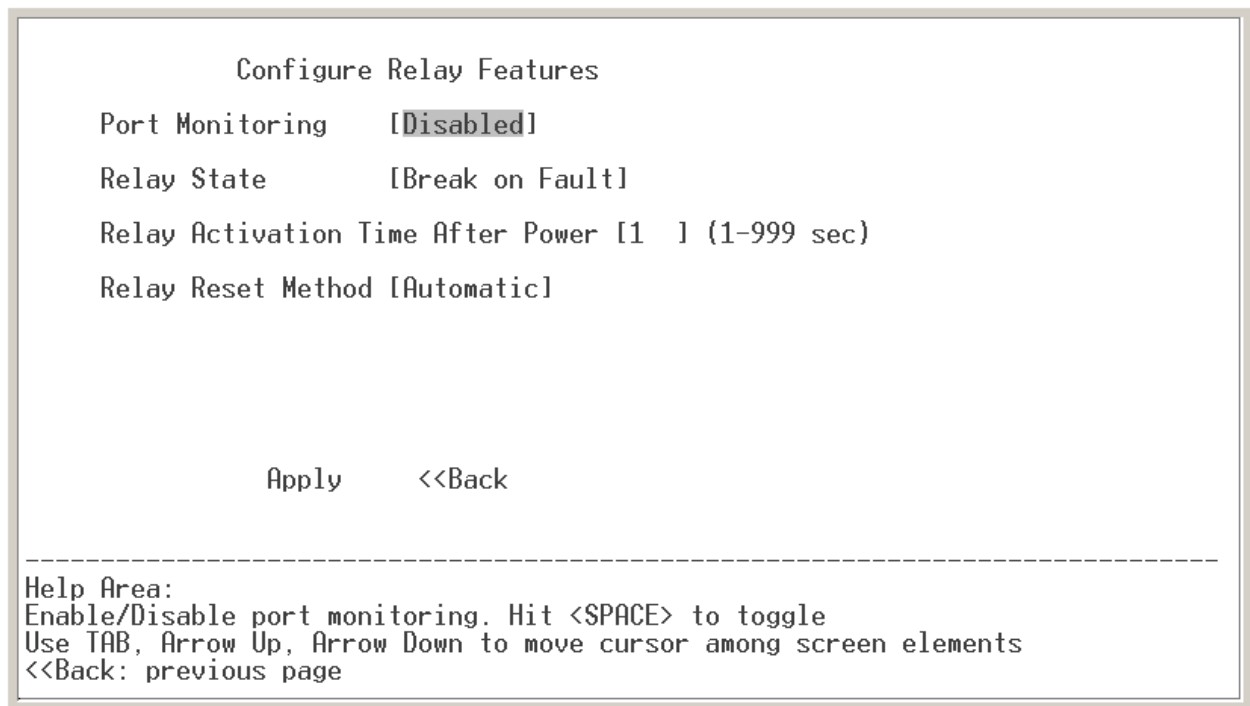


Figure 28 — Configure Relay Features

4.3.8.2 Configure Port Link Monitoring

As displayed for an 8-port switch in Figure 29, the user can monitor three conditions :

Ignore (the default) removes the port from link monitoring.

No Link reports a fault, (the relay to activate) if the link for the port has been lost.

Link Present reports a fault (the relay to activate) if a link has been detected on the port. This option is commonly used as a security feature to detect unauthorized connections to the switch.

```
Configure Port Link Monitoring
Port  Link Monitoring
  1   [No Link      ]
  2   [[Link Present]
  3   [[Ignore      ]
  4   [[Ignore      ]
  5   [[Ignore      ]
  6   [[Ignore      ]
  7   [[Ignore      ]
  8   [[Ignore      ]

Apply    <<Back

-----
Help Area:
Select port to monitor for No Link/Link Present/Ignore. Hit <SPACE> to toggle.
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 29 — Configure Port Link Monitoring

For 16- and 24-port switches, the applicable number of ports would be displayed in the above screen.

4.3.8.3 Display Faults

Figure 30 depicts an 8-port switch example of a **Display Faults** report. For switches of 16- and 24-ports, the applicable number of ports would be displayed. Except for the refresh option, this screen is read-only. The current fault report is displayed when the screen first appears. While on screen, this report is static; it will only be updated if the refresh option is selected.

```
Display Faults
Port  Link Monitoring
  1   [Fault  ]
  2   [No Fault]
  3   [No Fault]
  4   [No Fault]
  5   [No Fault]
  6   [No Fault]
  7   [No Fault]
  8   [No Fault]

Refresh  <<Back

-----
Help Area:
Press <ENTER> to refresh screen
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 30 — Display Faults

After relay activation is noted, the fault should be corrected. If **No Link** monitoring is in force, this will require restoration of a broken link or repair of the defective device to which the switch is connected. On the other hand, if **Link Present** monitoring is in use, removing the offending cable or end device will be required.

4.3.9 Configure Redundancy

Each managed switch from Contemporary Controls offers you a choice between the public standard protocols known as *Spanning Tree Protocol* or *Rapid Spanning Tree Protocol* and the proprietary redundancy protocol known as *RapidRing*[™]. By default, the screen of Figure 31 (**Configure Redundancy**) displays the public protocol RSTP and its basic parameter values for this switch.

To choose RapidRing as your redundancy scheme, position your cursor in the line which begins with the word **Redundancy** then press your keyboard space bar to toggle the selection from *STP/RSTP* to *RapidRing*. Once this is done, a new screen appears with options listed in Section 4.2.1.9.3.4 — but before adjusting these values, you should be familiar with all of the material discussed in Section 4.2.1.9.3.

4.3.9.1 Rapid Spanning Tree Protocol

Rapid Spanning Tree Protocol (RSTP) provides network path redundancy but without data loops that are prohibited. If more than one *active* path exists between two stations, a confused forwarding algorithm could transmit duplicate frames — one along each path.

RSTP constructs a tree of all RSTP-compliant switches in the network. To avoid loops, it forces each redundant path into an inactive state. If a segment is interrupted, an algorithm reconfigures the tree by quickly activating a normally unused link to substitute for the failed link.

By the exchange of messages, each switch in the tree collects information on all other switches. This information includes switch and port priorities, Media Access Control (MAC) addresses and path merit figures called “port costs”. This exchange results in the election of one switch to perform as the **root switch** (the *logical* centre of the tree) and also defines how ports are to be used on all other switches. The **root port** on a switch will send traffic to the root switch along the most efficient path. If the root port is disrupted, a **backup port** is activated as the substitute. A **designated port** provides the best path for root-bound traffic from outlying switches. If the designated port is disrupted, an **alternate port** is activated as the substitute.

If all RSTP-compliant switches in the network are enabled with default settings, the switch with the lowest MAC address is elected the root switch. But due to network traffic and architecture issues, the elected switch might not be the best to serve as the root device. You can manually force the switch of your choice to serve as the root device by increasing its priority so that the root-election algorithm chooses it as the root.

In general, you should configure your RSTP network so that the paths with the greatest bandwidth are those which support traffic for the root switch. For conveying root traffic, a fibre optic link would be preferred over a copper link and a 100 Mbps link would serve better than one operating at 10 Mbps. Also, the tree should consist of only of devices that are RSTP-compliant — non-compliant switches and hubs, if used at all, should only occupy the periphery of the tree because they will not forward the special messages needed for the construction and maintenance of the tree.

4.3.9.2 Configure Spanning Tree Protocol

The second field of Figure 31 (labelled **STP/RSTP**) allows you to toggle the protocol to be either *Disabled* (the default) or *Enabled*. The remaining value fields are for the:

Version	(explained in Section 4.2.1.9.2.1)
Hello Time	(explained in Section 4.2.1.9.2.2)
Priority	(explained in Section 4.2.1.9.2.3)
Max Age	(explained in Section 4.2.1.9.2.4)
Forward Delay	(explained in Section 4.2.1.9.2.5)

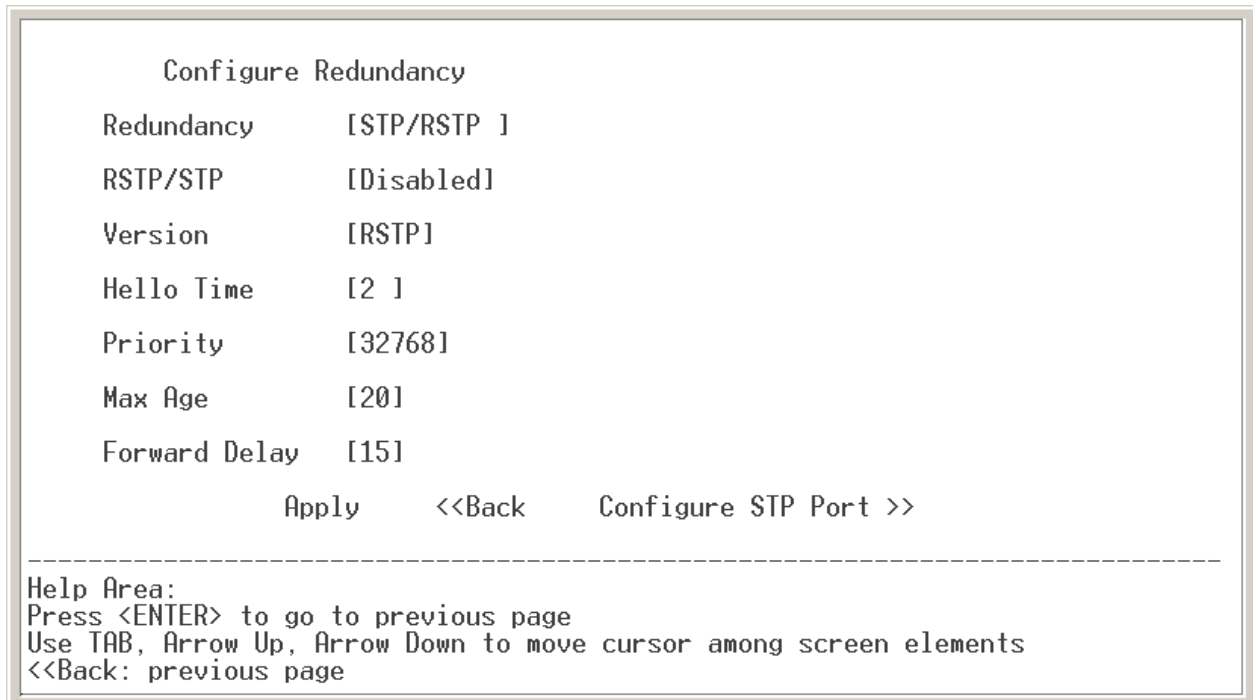


Figure 31 — Configure Redundancy

To configure the individual port RSTP parameters, select the option that is displayed as **Configure STP Port** — which is explained in Section 4.2.1.9.2.6.

4.3.9.2.1 **Version**

The original 1990 link-management specification in IEEE802.1D, Clause 8, described path redundancy via the Spanning Tree Protocol (STP). In 1999 STP was superseded by the Rapid Spanning Tree Protocol (RSTP) of IEEE802.1D, Clause 17. The RSTP interoperates with STP, but if RSTP-compliant switches are used in the same network with legacy STP-compliant switches, rapid reconfiguration may not be possible. RSTP is the default selection for managed switches from Contemporary Controls.

4.3.9.2.2 **Hello Time**

This is the interval at which the root device transmits a configuration message. The default value is 2 seconds, and can be set from 1–10 seconds. But if the result of the equation

$$(Max\ Age / 2) - 1$$

is less than 10, the maximum Hello Time will be the calculated value.

4.3.9.2.3 **Priority (of the entire switch)**

You can adjust the *switch* priority in steps of 4096 (the default is 32768) as follows:

0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, 61440.

4.3.9.2.4 **Max Age**

This is the maximum number of seconds a device waits to receive a configuration frame before attempting to reconfigure. The default value is 20 seconds, and it can be set from 6–40 seconds. However, actual minimum and maximum limits may be imposed by calculations based on the Hello Time and Forward Time as follows:

If the result of the equation

$$2 \times (Hello\ Time + 1)$$

is more than 6, the minimum Max Age will be the calculated value.

If the result of the equation

$$2 \times (Forward\ Time - 1)$$

is less than 40, the maximum Max Age will be the calculated value.

4.3.9.2.5 **Forward Delay**

This is the time a device will wait before changing states. Each device must receive topology information before it forwards frames and each port needs time to listen for any information that might force it to a discarding state. The default value is 15 seconds, and can be set from 4–30 seconds. But if the result of the equation

$$(Max\ Age / 2) + 1$$

is more than 4, the minimum Forward Delay will be the calculated value.

NOTE: The value limits for **Hello Time**, **Max Age** and **Forward Delay** will only be imposed if the RSTP/STP field has been *Enabled*.

4.3.9.2.5.1 Configure STP Port Parameters

Figure 32 displays the default settings for the **Configure STP Port Parameters** screen. This example is for an 8-port switch. A **Continue** option for accessing more ports would appear just above the dashed line for 16- and 24-port switches. The parameters are:

Path Cost	(explained in Section 4.2.1.9.2.6.1)
Priority	(explained in Section 4.2.1.9.2.6.2)
Admin Edge	(explained in Section 4.2.1.9.2.6.3)

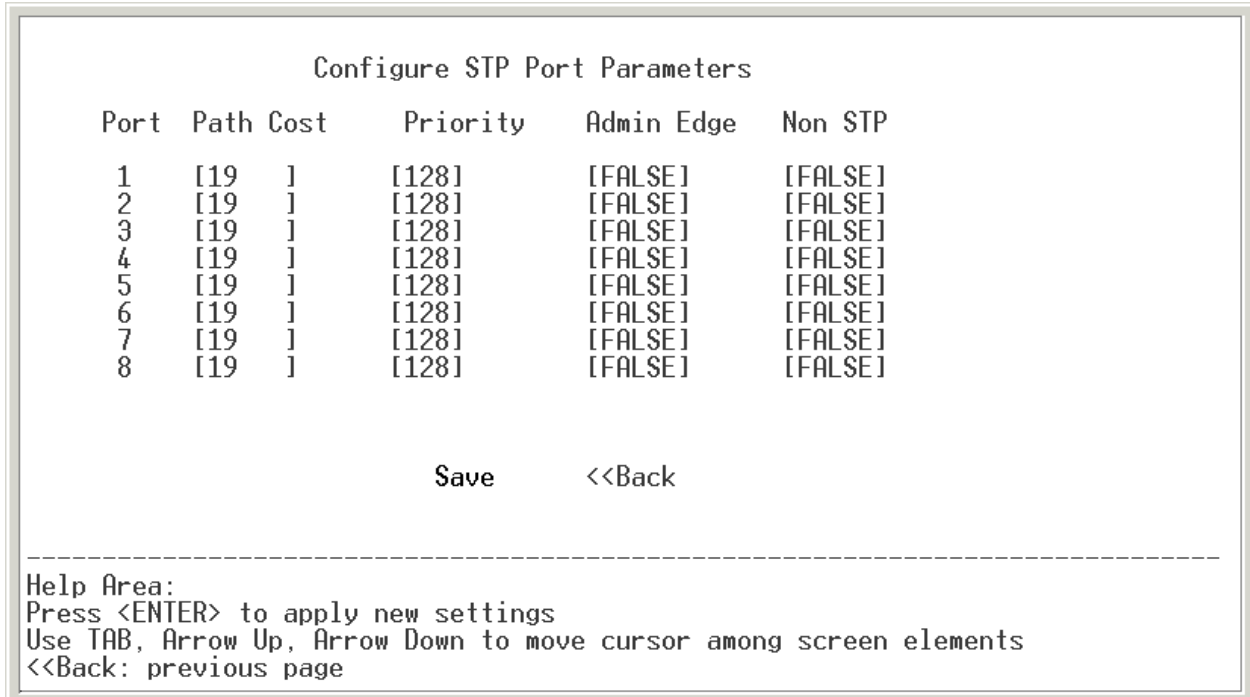


Figure 32 — Configure STP Port Parameters

4.3.9.2.5.2 Path Cost

In determining the most efficient path for conveying messages between the periphery of the tree and its root, one of the factors RSTP relies on is “path costs”. A typical link operates at either 10 or 100 Mbps and a port sending traffic to that link is assigned a “cost” derived from the link data rate. The default for the **Path Cost** field is 19 (the nominal RSTP port cost for a 100 Mbps link). Nominal values of port path costs and the *suggested* ranges through which these values might vary in most networks are listed in Table 1.

Data Rate	Cost Range	Cost Value
4 Mbps	100 – 1000	250
10 Mbps	50 – 600	100
16 Mbps	40 – 400	62
100 Mbps	10 – 60	19

Table 1 — Port Path Costs

Path Cost settings in Figure 32 can be set from 1–65535. When should you set a port path cost to a very high value? Although the associated link might operate at 10 Mbps, for example, the non-RSTP end of the link might have a very slow device such as a dial-up modem that could slow traffic drastically. In such a situation, you would likely want to raise this port path cost value to force RSTP to only use this path as a last resort.

4.3.9.2.5.3 Priority (of individual ports)

With the **Priority** field, you can modify the priorities of individual ports to affect RSTP path choices in the local vicinity of the switch. A lower value means a higher priority. Priority settings differ from path costs which are *cumulative* in calculating a total path from periphery to root. The **Priority** value only acts *locally* so you can force RSTP to favour a certain path emerging from the switch in question when two paths from the switch are otherwise equal. This field’s default value is 128 and its value can be toggled in 16 steps from 0–240 where each increment has a value of 16.

4.3.9.2.5.4 Admin Edge

This option (set to *FALSE* by default) can be set *TRUE* if the attached device falls outside the RSTP tree. Such a device could be an operator work station or a server. In this case, the port affected by the **Admin Edge** would be at the *edge* of RSTP *administration*. End nodes cannot cause loops, so they can pass directly into the *Forwarding* state. Setting this value to *TRUE* provides several benefits:

- gives faster RSTP algorithm solution (convergence)
- reduces flooding needed for rebuilding address tables during reconfiguration
- eliminates an RSTP reconfiguration if the port changes state
- improves other RSTP-related timeout issues

4.3.9.3 RapidRing™

4.3.9.3.1 Characteristics of RapidRing

RapidRing technology from Contemporary Controls provides high speed redundancy in Ethernet networks. It allows recovery in under 300 ms. If desired, the ring can consist of a mixture of 8-port, 16-port and 24-port switches.

RapidRing is wired in a simple ring structure using ports 7 and 8 of each 8-port switch (as shown in Figure 33) or ports 1 and 2 on each 16- or 24-port switch. Every link in the ring must connect the odd-numbered ring port (7 or 1) of one switch to the even-numbered ring port (8 or 2) of the adjacent switch. A properly constructed ring will never have a link between identical ports of adjacent switches. The ring ports can be wired with copper or fibre optic cable, depending on the model of switch.




Figure 33 — RapidRing

As shown in Figure 33, one switch must be selected as the “master” that activates the backup link if a ring failure occurs. The backup link must connect to the master through its **even-numbered** ring port which remains inactive during normal network activity. If a cable failure is detected, the master will activate its backup port to maintain communications.

NOTE: Upon enabling RapidRing, the switch will *automatically reboot*.

Fault relays and status LEDs work independent of RapidRing. If RapidRing is enabled *and* the ring ports are being monitored, the following applies: When a break occurs, each switch losing connectivity will flash its Status LED and activate its Fault Relay (and *regardless of port monitoring*, it will transmit a link-down SNMP trap). The break is thus located to the link between the two fault-reporting switches. Once the cable is repaired by the user, the fault relays will disengage and the status LEDs will glow solid to indicate the ring network is properly connected.

NOTE: A flashing Status LED and Fault Relay activation do not necessarily indicate a *ring* failure. The user may also be monitoring *non-ring* ports — in which case the failure report might indicate a non-ring issue. However, ring failure in a properly constructed circuit will be reported by *two ring* switches — a non-ring issue will not be so reported.

4.3.9.3.1 RapidRing and Other Management Features

As a rule, when ring ports are enabled for RapidRing use, they should not be involved with other management features. Specific issues are described below.

Port Parameters — If RapidRing is enabled, the “Configure Ports” options (Figure 6) are not available for the ring ports. That is, ring ports cannot be disabled nor can their configuration be changed.

Trunking — RapidRing and trunking *cannot* use the same ports. If a ring port exists as part of a trunk, the ring cannot be enabled. Either remove the ring port(s) from the trunk or disable the trunk that includes the ring port(s). Otherwise, the ring cannot be established. If a ring is enabled, its ports cannot be added to a trunk.

Mirroring — RapidRing ports may be mirrored like any other port, but neither of the ring ports should be designated as a Mirror Port (see Section 4.2.1.4).

VLANs — Ring ports must be included in the *same* VLAN or in none at all.

Multicast Filtering and Static Forwarding — RapidRing ports may be used for these functions, but the even-numbered ring port of the *ring master* is normally unavailable.

QoS — RapidRing ports do support QoS.

4.3.9.3.2 Multiple RapidRings

RapidRing will also support up to 100 interconnected rings (each having its own ID number) — allowing greater flexibility of wiring and network styles. Figure 34 shows two rings connected via two redundant network links — thus protecting against failure of a single cable within the trunk. (However, the two rings could be connected by just one link, if inter-ring redundancy were unneeded.) The two switches providing the inter-ring connection may be either master or slave but must use only non-ring ports for the trunk.

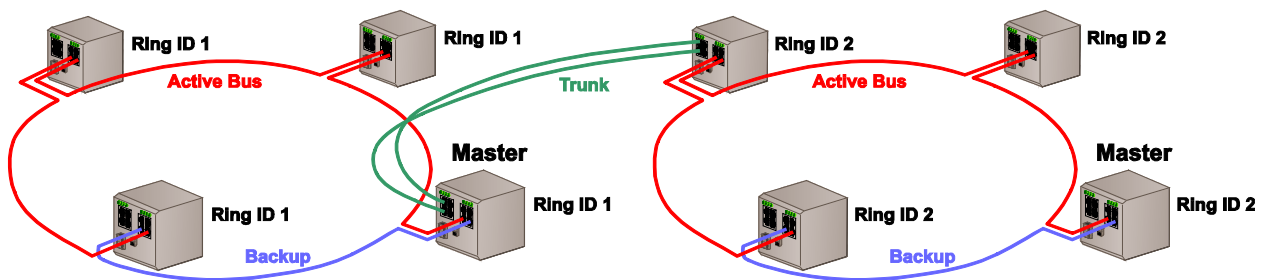


Figure 34 — Dual RapidRings

To setup the RapidRing via the console port settings, select **System Configuration**, then *Configure RapidRing* (described in Section 4.2.1.9.3.4). Also, if the two rings are to have redundant lines between them, configure a trunk between the two switches that connect the rings together. (See Section 4.2.1.3 for configuring a trunk).

4.3.9.3.2.1 Configure RapidRing

Each switch must be configured for ring operation. The console screen of Figure 35 displays the RapidRing *default* screen settings. There are three options to consider:

- **Ring Status** enables or disables RapidRing functionality. Every switch in a ring must have this option enabled. Otherwise, the presumed backup protection will not exist because a link failure might not be reported to the master.
- **Switch Mode** sets the master/slave status of the switch being configured. Only one master is defined per ring. Otherwise, the ring will not be established, some signal paths may not exist and messages could be lost.
- **Ring ID** must be a value from 1 – 100. The default value of 1 assumes that only one ring exists. If more rings are defined, each switch must be properly assigned to its Ring — otherwise, messages might be lost. When configuring multiple rings, all switches in a particular ring must have **matching** Ring ID values.

```
Configure Redundancy
Redundancy      [RapidRing]
Ring Status     [Disabled]
Switch Mode     [Slave   ]
Ring ID         [1]
Network Status: Ring Not Available

Refresh  Apply  <<Back

-----
Help Area:
Press <ENTER> to go to previous page
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 35 — Configure RapidRing

In addition to the three options described above, a read-only **Network Status** field reports the condition of the ring. The three following conditions are reported:

- **Ring Not Available** is displayed when RapidRing is not enabled.
- **Ring Incomplete** is displayed when RapidRing is enabled, but the master has invoked the backup link due to a primary ring failure.
- **Ring Complete** is displayed when RapidRing is enabled and all links are intact.

4.3.10 Configure Rate Control

4.3.10.1 Rate Control Overview

The *rate* and *type* of port traffic can be controlled to enhance network performance. This feature is deemed *Rate Control* and Figure 36 displays the menu screen for choosing **Ingress** or **Egress** port settings. With rate control, bandwidth allocation can be finely controlled.

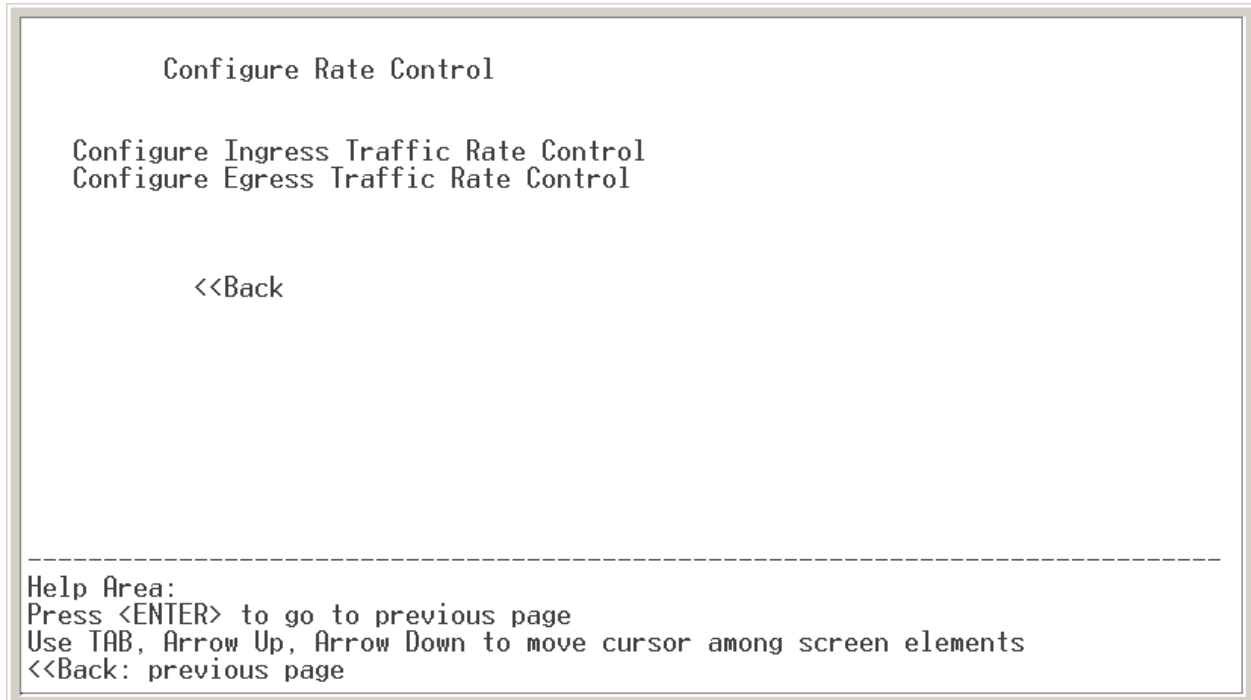


Figure 36— Configure Rate Control Menu

Rate control can be a useful feature for limiting communications from an unknown network. A common application of rate control is to prevent bandwidth “hogging” when interconnecting office and control networks.

The rate parameter is selectable for many different bit rates that can range from a high of 100 Mbps down to a low of 128 kbps. The default condition is for all messages to pass at 100 Mbps.

Selecting which frame types are to be used for rate control is accomplished on the *Configure Ingress Traffic Rate Control* screen described on the next page.

4.3.10.2 Configure Ingress Traffic Rate Control

Figure 37 shows a sample *Configure Ingress Traffic Rate Control* screen for a 24-port switch. An 8-port switch or a 16-port switch will have screens that display only the appropriate number of ports. Use your keyboard space bar to toggle through the available options while the cursor is positioned in the field being modified.

The frame types that you can control are *broadcast*, *multicast*, *unicast*, *destination lookup fail*, and *MAC control frame*. By placing *all types* under control, you can manage the total bandwidth of the port in question. Selecting only *broadcast* frames effectively creates a *broadcast storm control* that has a selectable maximum bandwidth setting. By controlling the *multicast* traffic, you can limit the extent to which a *group of devices* can consume bandwidth. On the other hand, if you deem group messaging more important, you can control the traffic of *unicast* messages. *Destination lookup fail* frames are merely unknown addresses — either never before encountered or those which have been aged out of the MAC lookup table. *MAC control frames* are special frames — the most common of which are PAUSE flow control frames.

In the example screen of Figure 37, rate control has been applied to broadcast and multicast frames as indicated by an “X” in their respective fields.

The selected Max Rate is the maximum bandwidth level for the types of messages selected. Types not selected will be allowed to use 100% of the port’s bandwidth.

NOTE: Port 1 in this example shows a Max Rate value set to “1/8”. This is 1/8 of a 1 Mbps data rate — which is 128 kbps, the minimum value.

```
Configure Ingress Traffic Rate Control

Control Type  [X] Broadcast          [X] Multicast
               [-] Unicast          [-] Destination Lookup Fail
               [-] MAC Control Frame

Port  Max Rate (Mbps)  Port  Max Rate (Mbps)  Port  Max Rate (Mbps)
1     [1/8]           9     [100]           17   [100]
2     [100]          10    [100]           18   [100]
3     [100]          11    [100]           19   [100]
4     [100]          12    [100]           20   [100]
5     [100]          13    [100]           21   [100]
6     [100]          14    [100]           22   [100]
7     [100]          15    [100]           23   [100]
8     [100]          16    [100]           24   [100]

Apply  <<Back

-----
Help Area:
Press <ENTER> to apply new settings
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 37 — Configure Ingress Traffic Rate Control

4.3.10.3 Configure Egress Traffic Rate Control

Figure 38 shows a sample *Configure Egress Traffic Rate Control* screen for a 24-port switch. As with the previous screen, the selected Max Rate is the maximum bandwidth level for the types of messages selected.

In the example screen of Figure 38, port 1 is selected and the available Max Rates are displayed in the help field near the bottom of the screen. The “1/8” of 1 Mbps rate is also available although this rate is not among those listed. Also, the management port only has two rate options: 100 Mbps and ¼ Mbps. Make your selection by pressing your keyboard’s space bar while the cursor is in the field to be modified.

```
Configure Egress Traffic Rate Control

Port  Max Rate (Mbps)  Port  Max Rate (Mbps)  Port  Max Rate (Mbps)
1     [100]              9     [100]             17    [100]
2     [100]             10    [100]             18    [100]
3     [100]             11    [100]             19    [100]
4     [100]             12    [100]             20    [100]
5     [100]             13    [100]             21    [100]
6     [100]             14    [100]             22    [100]
7     [100]             15    [100]             23    [100]
8     [100]             16    [100]             24    [100]
M     [100]

Apply    <<Back

-----
Help Area:
Max Rate,100/1/2/3/4/5/6/7/8/9/10/20/30/40/50/60/70/80/90(Mbps). <SPACE>=toggle
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 38 — Configure Egress Traffic Rate Control

4.3.11 Configure Port Security

Figure 39 illustrates an example of the *Configure Port Security* screen for an 8-port switch. Each port has its security disabled by default. If security is enabled for a port, no further MAC addresses are learned for that port and future transmissions through the port will only succeed if the source is listed in the address look-up table. Because static MAC addresses are not learned (aged out of the table), they are not affected by the applied security. Screens for 16- and 24-port switches would display the applicable number of ports for configuration.

```
Configure Port Security
Port Security
 1 [Disabled]
 2 [Disabled]
 3 [Disabled]
 4 [Disabled]
 5 [Disabled]
 6 [Disabled]
 7 [Disabled]
 8 [Disabled]

Apply <<Back Add Static MAC Address>>

-----
Help Area:
Press <ENTER> to apply new settings
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 39 — Configure Port Security

This is a useful feature when connecting to an unknown network — such as connecting the office network to the control network — and can be used to limit the office devices that can access the control network.

A convenient link to **Add Static MAC Addresses** is provided at the bottom of the screen. When activated, this link displays the same screen as shown in Figure 18.

4.3.12 Configure IGMP Snooping

Traditionally, IP messages are either unicast or broadcast, but multicasting can deliver messages to a select group of devices on the network. IGMP (Internet Group Multicast Protocol) is a session-layer protocol for defining membership in a multicast group. IGMP Group Destination Addresses (GDA) range from 224.0.0.0 to 239.255.255.255.

Managed switches from Contemporary Controls support IGMPv2 and can provide the querier function in the event no router exists in the network. Our switches implement the general query function in which members of all multicast groups report.

With IGMP Snooping, managed switches can recognize packets for multicast groups and direct them to only the destination ports which have received IGMP “join” messages from devices seeking membership in specific multicast groups. Figure 40 illustrates the *Configure IGMP Snooping* screen (for an 8-port switch). The controls are listed below.

IGMP Snooping State enables or disables IGMP Snooping.

IGMP Querier Function allows the switch to initiate a query to discover multicast group membership. Only one querier is allowed on a network at a time. If a query is transmitted by a device with a lower IP address than this switch, this switch relinquishes the role of querier to that device. If another LAN device is preferred for the query function, then this option can be disabled.

IGMP Forwarding Map specifies the ports that will forward IGMP join or leave messages to other querying devices. Disabling ports that do not connect to querying devices can improve bandwidth.

IGMP Query Interval (in seconds) specifies how often querying occurs. This setting is meaningless if the IGMP Query Function is disabled.

Multicast Filtering Age Out (in seconds) specifies how long the switch waits before deleting an entry from its multicast group list. If no member of a group responds within this time, the group is deemed inactive and removed from the list. This time should exceed the Query Interval or else the entry may be deleted before another query occurs. The switch can accommodate up to 100 multicast groups.

```
Configure IGMP Snooping

IGMP Snooping State   [Disabled]
IGMP Querier Function [Disabled]
IGMP Forwarding Map   [YYYYYYYY] (1-8)
Query Interval        [125 ] (30-9999 sec)
Multicast Age Out Timer[125 ] (30-9999 sec)

Apply    <<Back

-----
Help Area:
Press <ENTER> to go to previous page
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page.
```

Figure 40 — Configure IGMP Snooping

4.4 SNMP Configuration.

In a large network, a network management system (NMS) based on the Simple Network Management Protocol (SNMP) is often used to keep track of operations. SNMP, created in 1988, is the standard protocol for managing network devices. Over TCP/IP, SNMP usually uses UDP ports 161 (SNMP) and 162 (SNMP-traps).

An SNMP implementation involves three areas of functionality : the managed devices, the SNMP agents, and the NMS. SNMP agents reside in network devices where they use MIBs (information specific to the device) to interface the devices with the NMS — which then monitors and controls devices via these agents.

SNMP achieves device management via a very small command set described in 5.1.6. The switch can be managed via SNMP. From the **SNMP Configuration** menu (Figure 41), three submenus can be accessed :

Configure System Information	(explained in Section 4.2.2.1)
Configure SNMP Community	(explained in Section 4.2.2.2)
Configure Trap Receivers	(explained in Section 4.2.2.3)

For more information on SNMP support within the switch, see the Appendix.

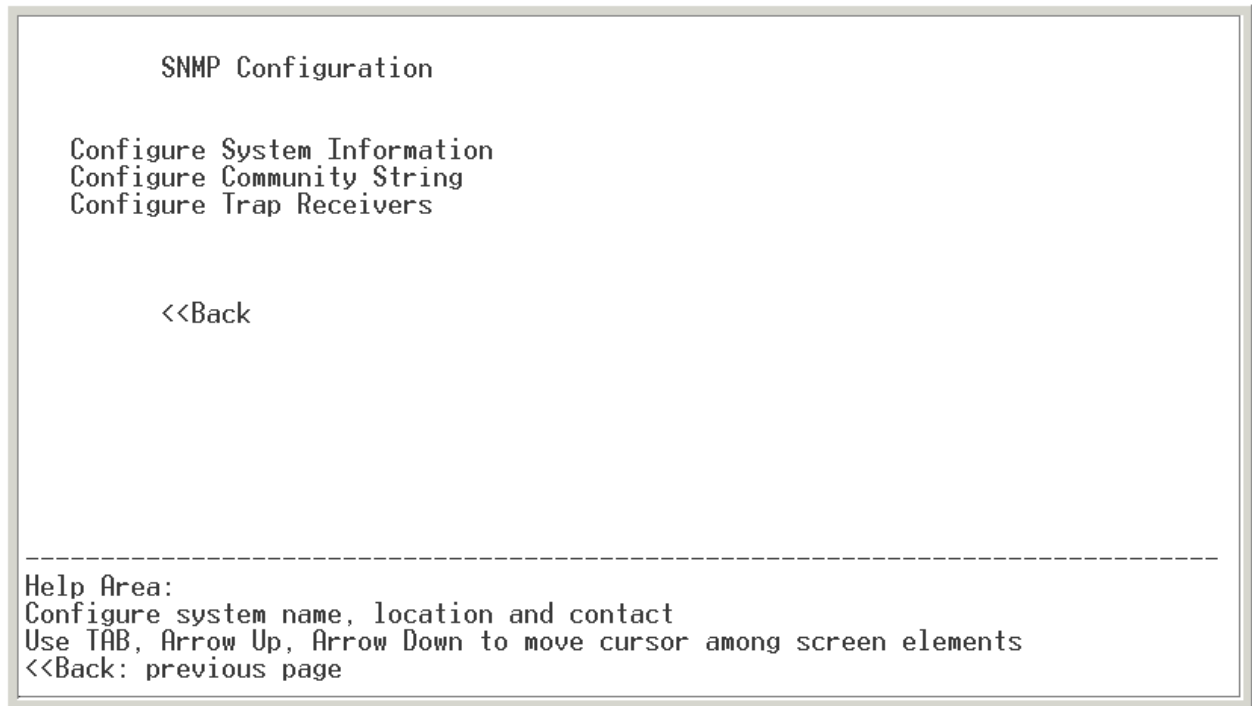


Figure 41 — SNMP Configuration

4.4.1 Configure System Information

Figure 42 shows the *Configure System Information* screen — although on the default screen you will see listed the actual System Name for your particular model of switch. This screen allows the setting of :

System Name	(1.3.6.1.2.1.1.5.0)
System Location	(1.3.6.1.2.1.1.6.0)
System Contact	(1.3.6.1.2.1.1.4.0)

These MIBs are among those listed in Section 5.1.1.1.

```
Configure System Information

System Name      [                                     ]
System Location [                                     ]
System Contact  [                                     ]

Apply    <<Back

-----
Help Area:
Enter system name
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 42 — Configure System Information

4.4.2 Configure SNMP Community

Managed devices are grouped into “communities” wherein every device has the same “community string” (aka “community name”) to be able to communicate via SNMP. This string assures authorized access to SNMP. Community strings can provide two types of access, read-only and read-write. Read-only access allows only *get* and *get-next* commands. Read-write access allows *get*, *get-next* and *set* commands.

The screen shown in Figure 43 allows the definition of up to four SNMP community names — each specifying either read-only access or read/write access. Before an access can be used, it must be set to *Valid*. Each community has three parameters to be configured :

Content is the string name (up to 10 characters) created by the user. It functions as a password to be used by any SNMP management software which accesses the switch.

Access is toggled by the user to be either RONLY (read-only) or RWRITE (read-write).

Status is toggled to be either *Valid* (string enabled) or *Invalid* (string disabled).

The example of Figure 43 has two community strings defined and valid : public (set for read-only access) and private (set for read-write access). (Several commonly-available SNMP manager applications use the terms *public* and *private* as default strings.)

```
Configure SNMP Community

Content      Access      Status
[public     ] [RONLY ] [Valid ]
[private    ] [RWRITE] [Valid ]
[           ] [RONLY ] [[Invalid]
[           ] [RONLY ] [[Invalid]

Apply      <<Back

-----
Help Area:
Enter SNMP community string
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 43 — Configure SNMP Community

4.4.3 Configure SNMP Trap Receivers

An SNMP Trap is a message that is transmitted when a trap event occurs. The menu in Figure 44 allows up to four trap receivers and each must be marked *Valid* for it to be used. Each valid trap receiver will receive a trap message upon a trap event occurring. The switch supports traps for :

link-up
link-down
authentication failure
cold start
warm start

The example of Figure 44 defines an **IP Address** for only one trap receiver — which will function because its status parameter has been toggled to *Valid*.

```
Configure SNMP Trap Receivers

IP Address      Status
[10.0.0.124    ] [Valid ]
[0.0.0.0       ] [Invalid]
[0.0.0.0       ] [Invalid]
[0.0.0.0       ] [Invalid]
Apply          <<Back

-----
Help Area:
Enter IP address of trap receiver
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 44 — Configure SNMP Trap Receivers

4.5 Performance Monitoring

The switch performance can be monitored via SNMP and console menus. The console menus support five main areas of monitoring. These are :

Monitor Port Traffic	(explained in Section 4.2.3.1)
Browse Address Table	(explained in Section 4.2.3.2)
Monitor Switch History	(explained in Section 4.2.3.3)
Monitor Switch Temperature	(explained in Section 4.2.3.4)
Monitor STP Port Status	(explained in Section 4.2.3.5)

```
Performance Monitoring

Monitor Port Traffic
Browse Address Table
Monitor Switch History
Monitor Switch Temperature
Monitor STP Port Status

<<Back

-----
Help Area:
Press <ENTER> to go to previous page
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 45 — Performance Monitoring

4.5.1 Monitor Port Traffic

Under the **Monitor Port Traffic** menu of Figure 46 there are two choices :

Port Packet Statistics (explained in Section 4.2.3.1.1)

Port Packet Error Statistics (explained in Section 4.2.3.1.2)

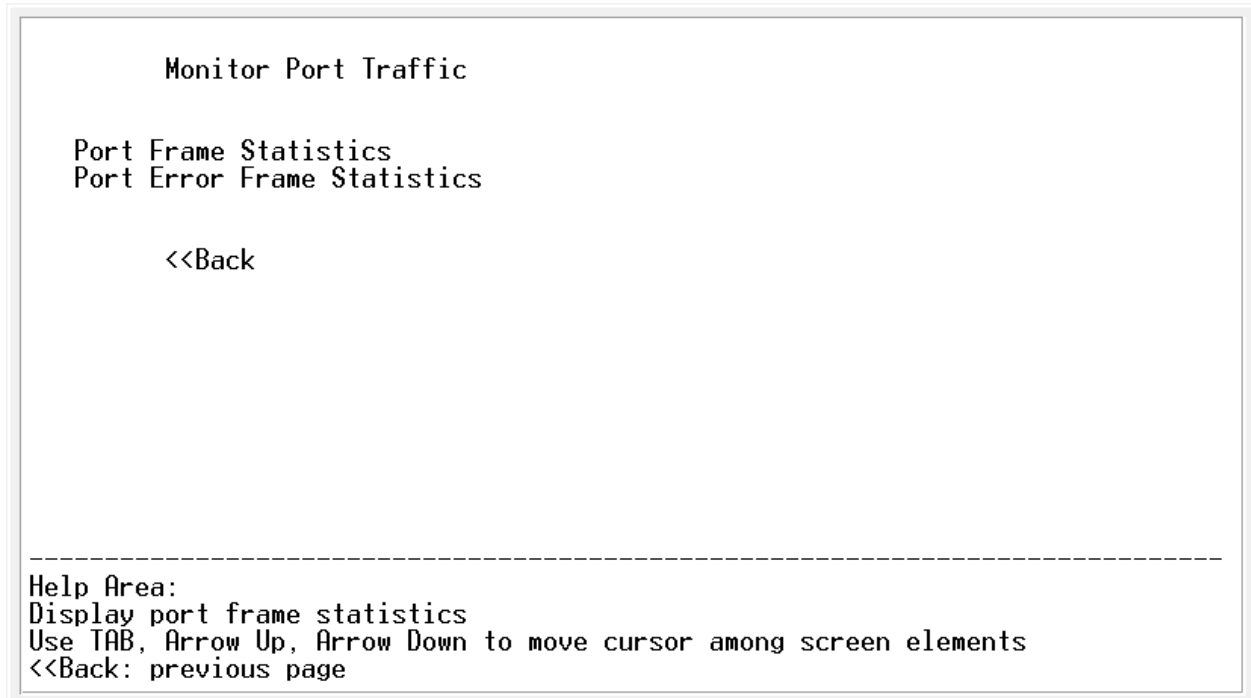


Figure 46 — Monitor Port Traffic

4.5.1.1 Port Frame Statistics

When monitoring port frame statistics, the non-error frame statistics for each port are displayed (Figure 47 is a sample for an 8-port switch). In this screen, the port number is first entered and then **Refresh** is selected to view the statistics for the entered port. These numbers will remain static until the **Refresh** option is again selected to provide an update. The displayed values are the total number of these events from when the switch was last powered-up, its IP address was redefined or its parameters were reset to their default values. Recycling power, redefining the IP address or resetting parameters to their default values will reset the **Port Frames Statistics** to zero.

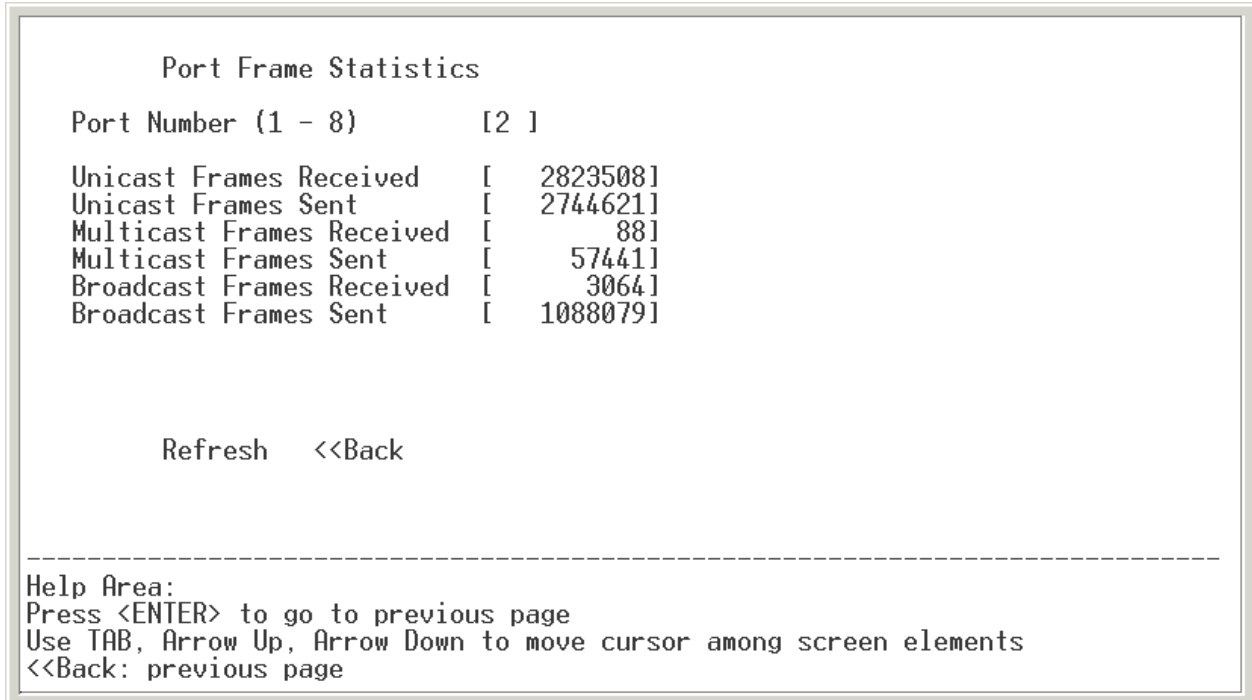


Figure 47 — Port Frame Statistics

4.5.1.2 Port Error Frames Statistics

Port Frame Error Statistics (Figure 48 is an example for an 8-port switch) are viewed by first entering the port number, then selecting the **Refresh** option. These values will remain static until updated with the **Refresh** option. The displayed values are the total number of these events from when the switch was last powered-up, its IP address was redefined or its parameters were reset to their default values. Recycling power, redefining the IP address or resetting parameters to their default values will reset the **Port Error Frames Statistics** to zero.

```
Port Error Frame Statistics

Port Number (1-8)      [ 2 ]

Dropped Frames        [          0]
Oversize Frames       [          0]
Undersize Frames      [          0]
Fragments             [          0]
Jabbers               [          0]
Collisions            [          0]
Deferred Transmission [          0]

Refresh  <<Back

-----
Help Area:
Press <ENTER> to go to previous page
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 48 — Port Error Frames Statistics

4.5.2 Browse Address Table

The **Browse Address Table** is displayed in Figure 49. The entire address table (up to 512 entries) can be displayed or a particular MAC address can be located. In the field named **Sort by**, use the spacebar to select the type of search you wish to have performed (*Sequence* or *MAC Address*) then select the **Find** menu option.

If multiple pages are required, the menu options **PreviousPage** and **NextPage** will be displayed. These allow the viewing of all located MAC addresses. The MAC addresses and the ports to which they are associated will also be displayed. The *VALID* status indicates the entry has not been aged out of the table. An *INVALID* status indicates the entry has been aged out of the table or deleted by the user and may be replaced when a new entry is added. The *STATIC* status indicates the entry is controlled by the management CPU and automatic learning and aging of the entry will not take place.

```

Browse Address Table

Sort by: [Sequence ]                               Find  <<Back

Total Entries in the Table (Max 512):  26

Number  MAC Address   Status  Port
001     0050DB00026F  VALID   01
002     00000000066F  VALID   01
003     000802558B5D  VALID   01
004     0050DAB7852F  VALID   05
005     00C002374254  VALID   01
006     00508BE7CEE0  VALID   02
007     00021761C8EE  VALID   01
008     00A0CC3BC04D  VALID   01
009     0050DB001330  STATIC  M
010     0050DB0000F3  VALID   08

                                           PreviousPage  NextPage
-----
Help Area:  Searching is done.
Start Searching in address table
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page. Note: M = management port

```

Figure 49 — Browse Address Table

4.5.3 Monitor Switch History

The **Monitor Switch History** screen in Figure 50 displays the number of seconds since the previous power cycle. The **Refresh** option will update the screen (the history counter is reset only as a result of a power cycle).

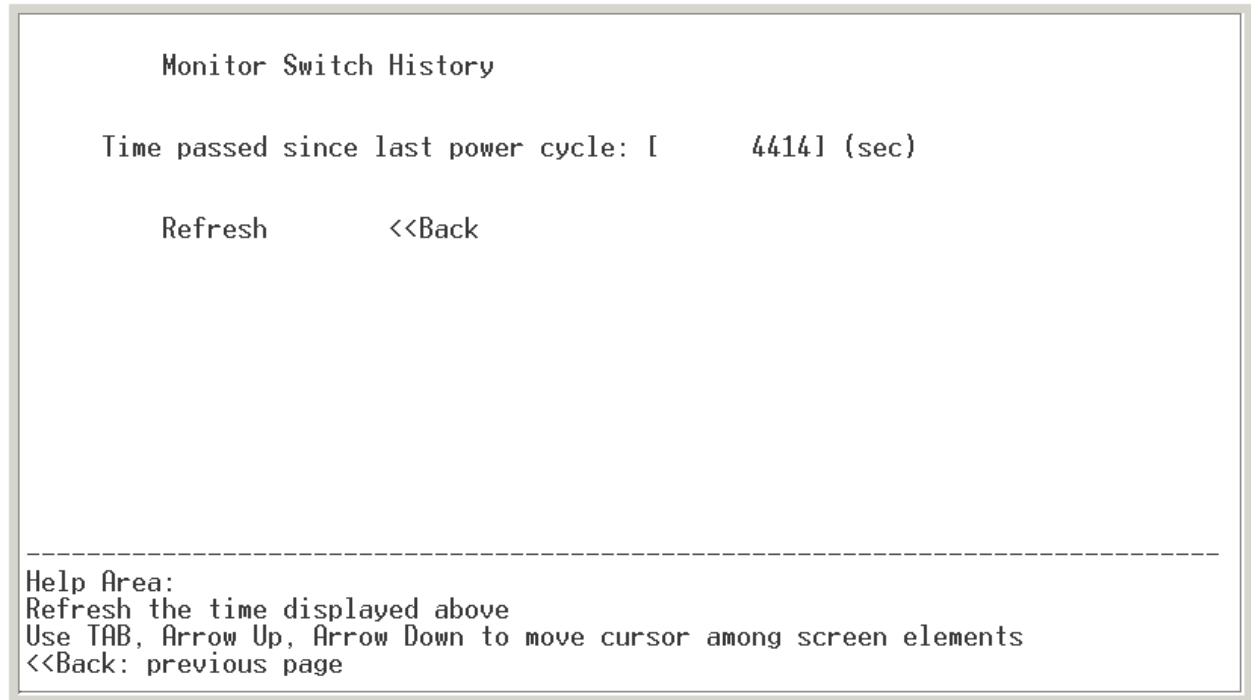


Figure 50 — Monitor Switch History

4.5.4 Monitor Switch Temperature

The **Monitor Switch Temperature** screen of Figure 51 displays the current internal temperature of the switch in degrees Celsius ($\pm 3^\circ$). Select the **Refresh** item will update the currently displayed value.

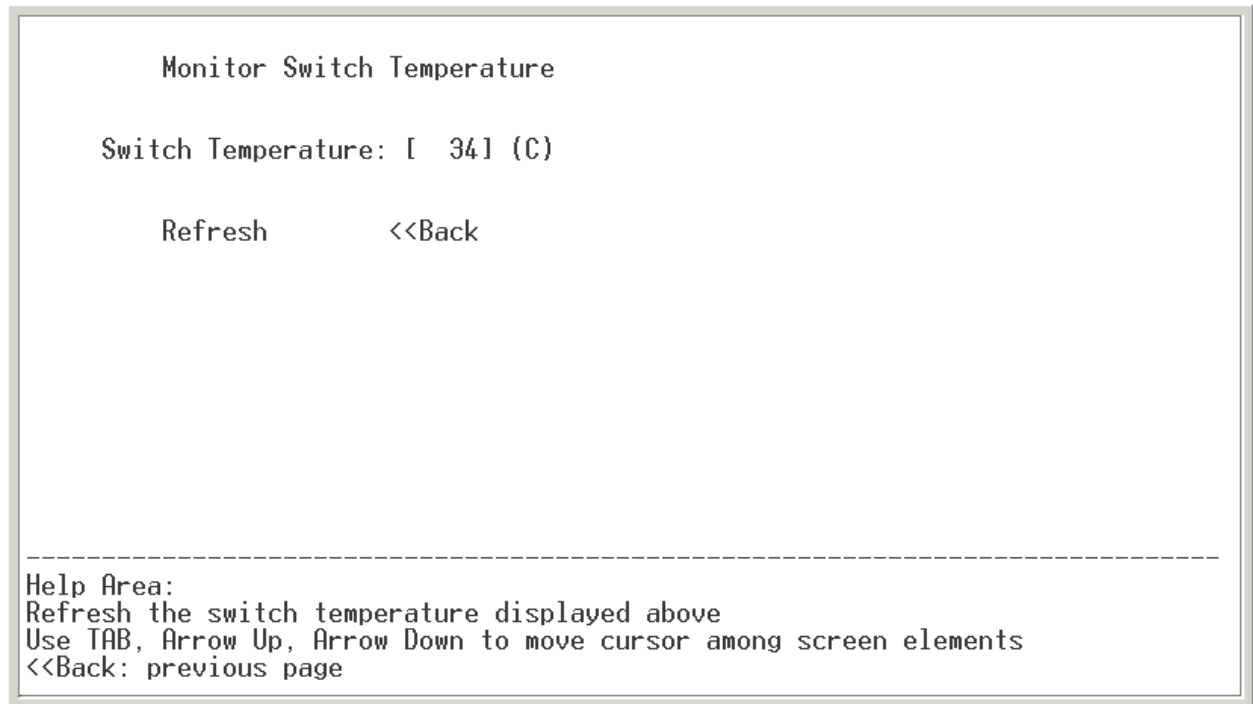


Figure 51 — Monitor Switch Temperature

4.5.5 Monitor STP Port Status

The screens of Figure 52 (where STP is enabled) and Figure 53 (where STP is disabled) report **STP State** and **Link Status** for each port. These examples show the first 8 ports of a 16- or 24-port device. The other ports would be accessed by choosing the *Continued* option just above the dashed line. The significance of the reported information is discussed on the next page.

```
Monitor STP Port Status

Port  STP State      Link Status      Speed/Duplex
 1  [Discarding]    [Down]          [      --      ]
 2  [Learning ]    [Up ]           [100Mbps Full Duplex]
 3  [Forwarding]   [Up ]           [100Mbps Full Duplex]
 4  [Discarding]   [Down]          [      --      ]
 5  [Discarding]   [Down]          [      --      ]
 6  [Discarding]   [Down]          [      --      ]
 7  [Discarding]   [Down]          [      --      ]
 8  [Discarding]   [Up ]           [100Mbps Full Duplex]

Root ID:  Self

Refresh  <<Back  Continued>>
-----
Help Area:
Press <ENTER> to go to previous page
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 52 — Monitor STP Port Status (STP Enabled)

```
Monitor STP Port Status

Port  STP State      Link Status      Speed/Duplex
 1  [Forwarding]   [Down]          [      --      ]
 2  [Forwarding]   [Up ]           [100Mbps Full Duplex]
 3  [Forwarding]   [Up ]           [100Mbps Full Duplex]
 4  [Forwarding]   [Down]          [      --      ]
 5  [Forwarding]   [Down]          [      --      ]
 6  [Forwarding]   [Down]          [      --      ]
 7  [Forwarding]   [Down]          [      --      ]
 8  [Forwarding]   [Up ]           [100Mbps Full Duplex]

Root ID:  Self

Refresh  <<Back  Continued>>
-----
Help Area:
Press <ENTER> to go to previous page
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
<<Back: previous page
```

Figure 53 — Monitor STP Port Status (STP Disabled)

STP State may report any of three states:

Forwarding indicates the RSTP port is actively participating in the tree — that is, it is not currently a backup or alternate port. The port forwards frames, and continues to learn new addresses. If RSTP is **disabled** as in Figure 53, each port is reported to be in the *Forwarding* state (whether the port is *Up* or *Down*) — but this has no meaning since the tree is disabled.

Learning is the brief port state (as it begins to learn addresses) before its forwarding delay expires. Actually, this state is rarely reported on screen, but could be if the user happens to refresh the monitor screen while a port is transitioning from *Discarding* to *Forwarding*. This *Learning* state will only occur as the tree is being restructured — as must be the case in the sample of Figure 52.

Discarding is reported if the port is not active in the tree. The port receives STP frames, but it does not forward frames. If the tree is stable and port states are consistent throughout the network, every root port and designated port will quickly transition through its *Learning* state to the *Forwarding* state. At the same time, all alternate and backup ports will stay in the *Discarding* state because they are not active in the tree.

Comments on Figures 52 and 53: A link may be reported as *Up* even if its associated port is not active in the tree. A look at Figure 52 reveals that although its link is *Up*, port 8 is *Discarding* — therefore port 8 is either a backup or an alternate port. Port 3 is listed as *Up* and *Forwarding* — thus, it is active in the tree and working properly. Port 2 is *Up*, but its state is momentarily *Learning* — so port 2 will soon become *Forwarding*.

Only RSTP port states (which superseded STP states) are shown in the monitor screen. For those wish to compare the port states of RSTP versus STP, refer to Table 2.

Status	STP State	RSTP State	Active In Tree?
Enabled	Blocking	Discarding	No
Enabled	Listening	Discarding	No
Enabled	Learning	Learning	Yes
Enabled	Forwarding	Forwarding	Yes
Disabled	Disabled	Discarding	No

Table 2 — STP & RSTP Port States Compared

The **Link Status** column of the Monitor STP Port Status screen simply reports if a port is *Up* or *Down*. It does this even if the RSTP/STP function is disabled.

Root ID is a read-only value of the MAC Address of the root switch. However, if the switch under consideration is the root switch, it reports as “Self”.

4.6 Upload/Download Settings

The switch settings can be stored to a PC and retrieved from a PC.

To **upload** the settings to the PC, select **Upload/Download Settings** from the *Managed Switch Main Menu* — like the one for model EISX8 displayed in Figure 54. In the next menu that appears, select *Upload Settings to PC*. Then, in **HyperTerminal** select *Transfer > Receive File*. Enter the name of the file and select the *Xmodem* protocol. Then press the *Receive* button.

NOTE: Settings uploaded to PC are read from the switch's non-volatile memory. Thus, to upload current settings, first store them to non-volatile memory (see Section 4.2) and then upload them.

To **download** a previously-uploaded settings file to the switch, select **Upload/Download Settings** from the *Managed Switch Main Menu*. In the next menu select *Download Settings from PC*. Then in **HyperTerminal** select *Transfer > Send File*. Enter the filename and select the *Xmodem* protocol. Then press the *Send* button. When the transfer is complete, recycle power on the switch to have the settings take effect.

```
EISX8 Managed Switch Main Menu

System Configuration
SNMP Configuration
Performance Monitoring
Username and Password
Save Settings to Non-Volatile Memory
Reset to Default Settings
Upload/Download Settings
Logout

Switch MAC Address: 00-50-DB-FF-40-12

-----
Help Area:
Configure switch management features
Use TAB, Arrow Up, Arrow Down to move cursor among screen elements
```

Figure 54 — Main Menu

5 Appendix

5.1 SNMP

The switch provides an SNMP interface for management of the device. The switch currently supports :

- RFC 1157 — SNMP protocol
- RFC 1215 — Traps for SNMP
- RFC 1213 — MIB-2
- RFC 1573 — MIB-2 Extension (IF-MIB)
- RFC 1493 — Bridge MIB
- RFC 1643 — Ethernet-like Interface MIB

The following MIBs are supported in the switch.

5.1.1 Managed Objects for TCP/IP Based Internet (MIB-II) — From RFC 1213

5.1.1.1 'System' group 1.3.6.1.2.1.1

oid = "1.3.6.1.2.1.1.1.0"

sysDescr: a textual description of the switch
Access: read-only

oid = "1.3.6.1.2.1.1.2.0"

sysObjectID: the vendor's authoritative identification
Access: read-only

oid = "1.3.6.1.2.1.1.3.0"

sysUpTime: the time since the last re-initialization
Access: read-only

oid = "1.3.6.1.2.1.1.4.0"

sysContact: the identification of the contact person
Access: read-write

oid = "1.3.6.1.2.1.1.5.0"

sysName: an administratively assigned name
Access: read-write

oid = "1.3.6.1.2.1.1.6.0"

sysLocation: the physical location of this node
Access: read-write

oid = "1.3.6.1.2.1.1.7.0"

sysServices: indicates the set of services that this switch primarily offers
Access: read-only

5.1.1.2 'Interfaces' group 1.3.6.1.2.1.2

oid = "1.3.6.1.2.1.2.1.0"

ifNumber: the number of network interfaces (regardless of their current state) present on this system.

Access: read-only

5.1.1.2.1 The Interfaces Table — 'ifTable' 1.3.6.1.2.1.2.2

oid = "1.3.6.1.2.1.2.2.1.1.ifIndex"

ifIndex: a unique value , greater than zero, for each interface.

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.2.ifIndex"

ifDescr: interface string contains information about the interface.

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.3.ifIndex"

ifType: interface type = 6 if Ethernet

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.4.ifIndex"

ifMtu: the size of the largest datagram that can be sent/received

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.5.ifIndex"

ifSpeed: interface speed = 100000000 for 100Base-TX

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.6.ifIndex"

ifPhysAddress: Ethernet (MAC) address (only used for designated management port on a switch)

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.7.ifIndex"

ifAdminStatus: the desired state of the interface

1 = up (ready to pass packets)

2 = down

3 = testing

Access: read-write

oid = "1.3.6.1.2.1.2.2.1.8.ifIndex"

ifOperStatus: the current operational state of the interface

1 = up (ready to pass packets)

2 = down

3 = testing

4 = unknown

5 = dormant

Access: read-only

oid = "1.3.6.1.2.1.2.2.1.9.ifIndex"
ifLastChange: the value of sysUpTime at the time the interface entered its current operational state
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.10.ifIndex"
ifInOctets: the total number of octets received on the interface, 32 bit.
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.11.ifIndex"
ifInUcastPkts: the number of unicast packets received, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.12.ifIndex"
ifInNUcastPkts: the number of non-unicast packets received, 32 bit.
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.13.ifIndex"
ifInDiscards: the number of inbound packets discarded with no error detected, 32 bit.
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.14.ifIndex"
ifInErrors: the number of inbound packets with errors, 32 bit
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.15.ifIndex"
ifInUnknowProtos: the number of packets received via the interface which were discarded because of unknown or unsupported protocol. Returns 0.
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.16.ifIndex"
ifOutOctets: the total number of packets transmitted, 32 bit.
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.17.ifIndex"
ifOutUcastPkts: the number of packets transmitted to a unicast address, 32 bit.
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.18.ifIndex"
ifOutNUcastPkts: the number of packets transmitted to a non-unicast address, 32 bit.
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.19.ifIndex"
ifOutDiscards: the number of outbound packets discarded with no error detected, 32 bit.
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.20.ifIndex"
ifOutErrors: the number of outbound packets with errors, 32 bit.
Access: read-only

oid = "1.3.6.1.2.1.2.2.1.22.ifIndex"
ifSpecific: a reference to MIB definition specific to the particular media being used to realize the interface.
Access: read-only

5.1.1.3 'IP' group 1.3.6.1.2.1.4

oid = "1.3.6.1.2.1.4.1.0"

ipForwarding: The indication of whether this switch is acting as an IP gateway in respect to the forwarding of datagrams received by, but not addressed to, this switch.

1 = forwarding

2 = not forwarding

Access: read-write

oid = "1.3.6.1.2.1.4.2.0"

ipDefault: The default value inserted into the Time-To-Live field of the IP header of datagrams originated at this switch, whenever a TTL value is not supplied by the transport layer protocol.

Access: read-only

oid = "1.3.6.1.2.1.4.3.0"

ipInReceives: The total number of input datagrams received from interfaces, including those received in error.

Access: read-only

oid = "1.3.6.1.2.1.4.9.0"

ipInDelivers: The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).

Access: read-only

oid = "1.3.6.1.2.1.4.10.0"

ipOutRequests: The total number of IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission.

Access: read-only

oid = "1.3.6.1.2.1.4.15.0"

ipReasmOKs: The number of IP datagrams successfully re-assembled."

Access: read-only

oid = "1.3.6.1.2.1.4.17.0"

ipFragOKs: The number of IP datagrams that have been successfully fragmented at this switch."

Access: read-only

5.1.1.3.1 The IP Address Table — 'ipAddrTable' 1.3.6.1.2.1.4.20

When IP address is used as input, its format should be 4 decimal fields.

oid = "1.3.6.1.2.1.4.20.1.1.<ipAdEntAddr>"

ipAdEntAddr: the IP address

Access: read-only

oid = "1.3.6.1.2.1.4.20.1.2.< ipAdEntAddr>"

ipAdEntIfIndex: physical port number associated with this particular subnet by IP address

Access: read-only

oid = "1.3.6.1.2.1.4.20.1.3.< ipAdEntAddr >"

ipAdEntNetMask: subnet mask associated with this particular subnet by IP address

Access: read-only

oid = "1.3.6.1.2.1.4.20.1.4.< ipAdEntAddr >"

ipAdEntBcastAddr: the value of the least-significant bit in the IP broadcast address. = 1 for Internet standard all-ones broadcast address.

Access: read-only

oid = "1.3.6.1.2.1.4.20.1.5.< ipAdEntAddr >"

ipAdEntReasmMaxSize: the size of the largest IP datagram which this switch can re-assemble from incoming IP fragmented.

Access: read-only

5.1.1.4 'ICMP' group 1.3.6.1.2.1.5

oid = "1.3.6.1.2.1.5.1.0"

icmpInMsgs: The total number of ICMP messages which the switch received.

Access: read-only

oid = "1.3.6.1.2.1.5.2.0"

icmpInErrors: The number of ICMP messages which the switch received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).

Access: read-only

oid = "1.3.6.1.2.1.5.8.0"

icmpInEchos: The number of ICMP Echo (request) messages received.

Access: read-only

oid = "1.3.6.1.2.1.5.9.0"

icmpInEchoReps: The number of ICMP Echo Reply messages received.

Access: read-only

oid = "1.3.6.1.2.1.5.14.0"

icmpOutMsgs: The total number of ICMP messages which this switch attempted to send.

Access: read-only

oid = "1.3.6.1.2.1.5.15.0"

icmpOutErrors: The number of ICMP messages which this switch did not send due to problems discovered within ICMP

Access: read-only

5.1.1.5 'TCP' group 1.3.6.1.2.1.6

oid = "1.3.6.1.2.1.6.10.0"

tcpInSegs: The total number of segments received, including those received in error.

Access: read-only

oid = "1.3.6.1.2.1.6.11.0"

tcpOutSegs: The total number of segments sent, including those on current connections but excluding those containing only retransmitted octets.

Access: read-only

oid = "1.3.6.1.2.1.6.12.0"

tcpRetransSegs: The total number of segments retransmitted.

Access: read-only

5.1.1.6 'UDP' group 1.3.6.1.2.1.7

5.1.1.6.1 The UDP Listener Table — 'udpTable' 1.3.6.1.2.1.7.5

When an IP address is used as input, its format should be 4 decimal fields.

oid = "1.3.6.1.2.1.7.5.1.1.< udpLocalAddress >.< udpLocalPort >"

udpLocalAddress: The local IP address for this UDP listener.

Access: read-only

oid = "1.3.6.1.2.1.7.5.1.2.<Local IP address as 4 decimal fields>.<LocalPort>"

udpLocalPort: The local port number for this UDP listener.

Access: read-only

5.1.1.7 'Transmission' group 1.3.6.1.2.1.10

Based on the transmission media underlying each interface on a system, the corresponding portion of the Transmission group is mandatory for that system.

In switches, all interfaces are Ethernet-based. Therefore, the Ethernet-like interface is used.

5.1.1.8 'SNMP' group 1.3.6.1.2.1.11

oid = "1.3.6.1.2.1.11.1.0"

snmplnPkts: The total number of Messages delivered to the SNMP switch from the transport service.

Access: read-only

oid = "1.3.6.1.2.1.11.2.0"

snmpOutPkts: The total number of SNMP Messages which were passed from the SNMP protocol switch to the transport service.

Access: read-only

oid = "1.3.6.1.2.1.11.3.0"

snmplnBadVersions: The total number of SNMP Messages which were delivered to the SNMP protocol switch and were for an unsupported SNMP version.

Access: read-only

oid = "1.3.6.1.2.1.11.4.0"

snmplnBadCommunityNames: The total number of SNMP Messages delivered to the SNMP protocol switch which used a SNMP community name not known to said switch.

Access: read-only

oid = "1.3.6.1.2.1.11.5.0"

snmplnBadCommunityUses: The total number of SNMP messages delivered to the SNMP protocol switch which represented an SNMP operation which was not allowed by the SNMP community named in the message.

Access: read-only

oid = "1.3.6.1.2.1.11.6.0"

snmplnASNParseErrs: The total number of ASN.1 or BER errors encountered by the SNMP protocol switch when decoding received SNMP Messages.

Access: read-only

oid = "1.3.6.1.2.1.11.8.0"

snmplnTooBig: The total number of SNMP PDUs which were delivered to the SNMP protocol switch and for which the value of the error-status field is 'tooBig'

Access: read-only

oid = "1.3.6.1.2.1.11.9.0"
snmplnNoSuchNames: The total number of SNMP PDUs which were delivered to the SNMP protocol switch and for which the value of the error-status field is 'noSuchName'.
Access: read-only

oid = "1.3.6.1.2.1.11.10.0"
snmplnBadValues: The total number of SNMP PDUs which were delivered to the SNMP protocol switch and for which the value of the error-status field is 'badValue'.
Access: read-only

oid = "1.3.6.1.2.1.11.11.0"
snmplnReadOnlys: The total number valid SNMP PDUs which were delivered to the SNMP protocol switch and for which the value of the error-status field is 'readOnly'.
Access: read-only

oid = "1.3.6.1.2.1.11.12.0"
snmplnGenErrs: The total number of SNMP PDUs which were delivered to the SNMP protocol switch and for which the value of the error-status field is 'genErr'.
Access: read-only

oid = "1.3.6.1.2.1.11.13.0"
snmplnTotalReqVars: The total number of MIB objects which have been retrieved successfully by the SNMP protocol switch as the result of receiving valid SNMP Get-Request and Get-Next PDUs.
Access: read-only

oid = "1.3.6.1.2.1.11.14.0"
snmplnTotalSetVars: The total number of MIB objects which have been altered successfully by the SNMP protocol switch as the result of receiving valid SNMP Set-Request PDUs.
Access: read-only

oid = "1.3.6.1.2.1.11.15.0"
snmplnGetRequests: The total number of SNMP Get-Request PDUs which have been accepted and processed by the SNMP protocol switch.
Access: read-only

oid = "1.3.6.1.2.1.11.16.0"
snmplnGetNexts: The total number of SNMP Get-Next PDUs which have been accepted and processed by the SNMP protocol switch.
Access: read-only

oid = "1.3.6.1.2.1.11.17.0"
snmplnSetRequests: The total number of SNMP Set-Request PDUs which have been accepted and processed by the SNMP protocol switch.
Access: read-only

oid = "1.3.6.1.2.1.11.18.0"
snmplnGetResponses: The total number of SNMP Get-Response PDUs which have been accepted and processed by the SNMP protocol switch.
Access: read-only

oid = "1.3.6.1.2.1.11.19.0"
snmplnTraps: The total number of SNMP Trap PDUs which have been accepted and processed by the SNMP protocol switch.
Access: read-only

oid = "1.3.6.1.2.1.11.20.0"
snmpOutTooBig: The total number of SNMP PDUs which were generated by the SNMP protocol switch and for which the value of the error-status field is 'tooBig.'
Access: read-only

oid = "1.3.6.1.2.1.11.21.0"
snmpOutNoSuchNames: The total number of SNMP PDUs which were generated by the SNMP protocol switch and for which the value of the error-status is 'noSuchName'.
Access: read-only

oid = "1.3.6.1.2.1.11.22.0"
snmpOutBadValues: The total number of SNMP PDUs which were generated by the SNMP protocol switch and for which the value of the error-status field is 'badValue'.
Access: read-only

oid = "1.3.6.1.2.1.11.24.0"
snmpOutGenErrs: The total number of SNMP PDUs which were generated by the SNMP protocol switch and for which the value of the error-status field is 'genErr'.
Access: read-only

oid = "1.3.6.1.2.1.11.25.0"
snmpOutGetRequests: The total number of SNMP Get-Request PDUs which have been generated by the SNMP protocol switch.
Access: read-only

oid = "1.3.6.1.2.1.11.26.0"
snmpOutGetNexts: The total number of SNMP Get-Next PDUs which have been generated by the SNMP protocol switch.
Access: read-only

oid = "1.3.6.1.2.1.11.27.0"
snmpOutSetRequests: The total number of SNMP Set-Request PDUs which have been generated by the SNMP protocol switch.
Access: read-only

oid = "1.3.6.1.2.1.11.28.0"
snmpOutGetResponses: The total number of SNMP Get-Response PDUs which have been generated by the SNMP protocol switch.
Access: read-only

oid = "1.3.6.1.2.1.11.29.0"
snmpOutTraps: The total number of SNMP Trap PDUs which have been generated by the SNMP protocol switch.
Access: read-only

oid = "1.3.6.1.2.1.11.30.0"
snmpEnableAuthenTraps: Indicates whether the SNMP agent process is permitted to generate authentication-failure traps.
1 = enabled
2 = disabled
Access: read-write

5.1.2 Managed Objects for Bridges — From RFC 1493

Bridge MIB — 'dot1dBridge' 1.3.6.1.2.1.17

5.1.2.1 'dot1dBase' group 1.3.6.1.2.1.17.1

oid = "1.3.6.1.2.1.17.1.1.0"

dot1dBaseBridgeAddress: MAC address used by this bridge
Access: read-only

oid = "1.3.6.1.2.1.17.1.2.0"

dot1dBaseNumPorts: number of ports controlled by this bridging switch
Access: read-only

oid = "1.3.6.1.2.1.17.1.3.0"

dot1dBaseType: indicates what type of bridging this bridge can perform.
1 = unknown
2 = transparent-only
3 = sourceroute-only
Access: read-only

5.1.2.1.1 'dot1dBasePortTable' 1.3.6.1.2.1.17.1.4

oid = "1.3.6.1.2.1.17.1.4.1.1.port"

dot1dBasePort: the port number of the port for which this entry contains bridge management information .
Access: read-only

oid = "1.3.6.1.2.1.17.1.4.1.2.port"

dot1dBasePortIfIndex: the value of instance of ifIndex object defined in Interface group of MIB-2 for the interface corresponding to this port.
Access: read-only

oid = "1.3.6.1.2.1.17.1.4.1.3.port"

dot1dBasePortCircuit: for a port which has the same value of dot1BasePortIfIndex as another port on the same bridge, this object contains the name of an object instance unique to this port. For a port which has a unique value of dot1dBasePortIfIndex, this object can have the value {0}.
Access: read-only

oid = "1.3.6.1.2.1.17.1.4.1.4.port"

dot1dBasePortDelayExceededDiscards: the number of frames discarded by this port due to excessive transmit delay through the bridge. Returns 0.
Access: read-only

oid = "1.3.6.1.2.1.17.1.4.1.5.port"

dot1dBasePortMtuExceededDiscards: the number of frames discarded by this port due to an excessive size.
Access: read-only

5.1.2.2 'dot1dTp' group 1.3.6.1.2.1.17.4

oid = "1.3.6.1.2.1.17.4.2.0"

dot1dTpAgingTime: the timeout period in seconds for aging out dynamically learned forwarding information.
Access: read-write

5.1.2.3 *'dot1dTpFdbTable'* 1.3.6.1.2.1.17.4.3

oid = "1.3.6.1.2.1.17.4.3.1.1.<MAC address as 6 decimal fields>"

dot1dTpFdbAddress: a unicast MAC address for which the bridge has forwarding and/or filtering information.

Access: read-only

oid = "1.3.6.1.2.1.17.4.3.1.2.<MAC address as 6 decimal fields>"

dot1dTpFdbPort: port number where this MAC has been 'learned' and stored in the switch lookup table.

Access: read-only

oid = "1.3.6.1.2.1.17.4.3.1.3. <MAC address as 6 decimal fields>"

dot1dTpFdbStatus: the status of this entry.

1 = other

2 = invalid

3 = learned

4 = self

5 = mgmt

Access: read-only

5.1.2.4 *'dot1dTpPortTable'* 1.3.6.1.2.1.17.4.4

oid = "1.3.6.1.2.1.17.4.4.1.1.port"

dot1dTpPort: the port number of the port for which this entry contains transparent bridging management information.

Access: read-only

oid = "1.3.6.1.2.1.17.4.4.1.2.port"

dot1dTpPortMaxInfo: the maximum size of the INFO (non-MAC) field that this port will receive or transmit.

Access: read-only

oid = "1.3.6.1.2.1.17.4.4.1.3.port"

dot1dTpPortInFrames: the number of frames received by this port.

Access: read-only

oid = "1.3.6.1.2.1.17.4.4.1.4.port"

dot1dTpPortOutFrames: the number of frames transmitted by this port.

Access: read-only

oid = "1.3.6.1.2.1.17.4.4.1.5.port"

dot1dTpPortInDiscards: the number of valid frames received which were discarded by the forwarding process.

Access: read-only

5.1.3 Managed Objects for Ethernet-like Interface Types — From RFC 1643

Ethernet-like Interface MIB — ‘dot3’ 1.3.6.1.2.1.10.7

5.1.3.1 Ethernet-like Statistics Group — ‘dot3StatsTable’ 1.3.6.1.2.1.10.7.2

oid = “1.3.6.1.2.1.10.7.2.1.1.dot3StatsIndex”

dot3StatsIndex: an index that identifies an interface, same value as ifIndex.

Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.2.dot3StatsIndex”

dot3StatsAlignmentErrors: the number of frames received on the interface that are not an integral number of octets in length and do not pass the FCS check. This count is incremented when the alignmentError status is returned by the MAC service to the LLC.

Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.3.dot3StatsIndex”

dot3StatsFCSErrors: the number of frames received on the interface that are not an integral number of octets in length and do not pass the FCS check. This count is incremented when the frameCheckError status is returned by the MAC service to the LLC.

Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.4.dot3StatsIndex”

dot3StatsSingleCollisionFrames: the number of successfully transmitted frame on the interface for which transmission is inhibited by one collision.

Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.5.dot3StatsIndex”

dot3StatsMultipleCollisionFrames: the number of successfully transmitted frame on the interface for which transmission is inhibited by more than one collision.

Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.6.dot3StatsIndex”

dot3StatsSQETestErrors: the number of times that the SQE TEST ERROR message is generated by the PLS sublayer for this interface. Returns 0.

Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.7.dot3StatsIndex”

dot3StatsDeferredTransmissions: the number of frames for which the first transmission attempt on the interface is delayed because the medium is busy.

Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.8.dot3StatsIndex”

dot3StatsLateCollisions: the number of times that a collisions is detected on the interface later than 512 bit-times into the transmission of a packet.

Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.9.dot3StatsIndex”

dot3StatsExcessiveCollisions: the number of frames for which transmission on the interface fails due to excessive collisions.

Access: read-only

oid = “1.3.6.1.2.1.10.7.2.1.10.dot3StatsIndex”

dot3StatsInternalMacTransmitErrors: the number of frames for which transmission on the interface fails due to an internal MAC sublayer transmit error.

Access: read-only

oid = "1.3.6.1.2.1.10.7.2.1.11.dot3StatsIndex"

dot3StatsCarrierSenseErrors: the number of frames that the carrier sense condition was lost or never asserted when attempting to transmit a frame on this interface.

Returns 0.

Access: read-only

oid = "1.3.6.1.2.1.10.7.2.1.13.dot3StatsIndex"

dot3StatsFrameTooLong: the number of frames received on a particular interface that exceed the maximum permitted frame size.

Access: read-only

oid = "1.3.6.1.2.1.10.7.2.1.16.dot3StatsIndex"

dot3StatsInternalMacReceiveErrors: the number of frames for which reception on the interface fails due to an internal MAC sublayer transmit error.

Access: read-only

oid = "1.3.6.1.2.1.10.7.2.1.17.dot3StatsIndex"

dot3StatsEthernetChipSet: this object contains an OBJECT IDENTIFIER which identifies the chipset used to realize the interface.

Access: read-only

5.1.4 Evolution of the Interface Group of MIB-II — From RFC 1573

MIBs for generic objects for network Interface sub-layers — ‘ifMIB’ 1.3.6.1.2.1.31

- *MIB Objects — ‘ifMIBObjects’ 1.3.6.1.2.1.31.1*
- *Extension to the Interface Table — ‘ifXTable’ 1.3.6.1.2.1.31.1.1*

oid = “1.3.6.1.2.1.31.1.1.1.ifIndex”

ifName: the textual name of the interface.
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.2.ifIndex”

ifInMulticastPkts: the number of multicast packets received, 32 bit.
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.3.ifIndex”

ifInBroadcastPkts: the number of broadcast packets received, 32 bit.
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.4.ifIndex”

ifOutMulticastPkts: the number of multicast packets transmitted, 32 bit.
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.5.ifIndex”

ifOutBroadcastPkts: the number of broadcast packets transmitted, 32 bit.
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.6.ifIndex”

ifHCInOctets: the total number of octets received on the interface, 64 bit.
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.10.ifIndex”

ifHCOutOctets: the total number of octets transmitted by the interface, 64 bit.
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.14.ifIndex”

ifLinkUpDownTrapEnable: indicates whether or not linkup/linkDown traps should be generated for this interface.
1 = enabled
2 = disabled
Access: read-write

oid = “1.3.6.1.2.1.31.1.1.1.15.ifIndex”

ifHighSpeed: an estimate of the interface current bandwidth in units of 1M bits/sec.
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.16.ifIndex”

ifPromiscuousMode: indicates whether this interface only accepts packets/frames addressed to this station (false) or accepts all packets/frames transmitted on the media (true). The value does not affect broadcast and multicast packets/frames.
1 = true
2 = false
Access: read-only

oid = “1.3.6.1.2.1.31.1.1.1.17.ifIndex”

ifConnectorPresent: indicates whether the interface has a physical connector.
1 = true
2 = false
Access: read-only

5.1.5 Private Managed Objects

MIBs for Contemporary Controls — 1.3.6.1.4.1.17384

- **Switch Family — 1.3.6.1.4.1.17384.1**
- **Switch Series — 1.3.6.1.4.1.17384.1.1**

oid = "1.3.6.1.4.1.17384.1.1.1.0"

temperature: indicates the internal temperature of the switch in degrees Celsius.

Type: string

Access: read-only

5.1.5.1 Relay Group – 1.3.6.1.4.1.17384.1.1.2

oid = "1.3.6.1.4.1.17384.1.1.2.1.0"

switchFaultStatus: the fault status of this switch.

Type: integer

1 = fault has occurred

2 = no fault has occurred

Access: read-only

5.1.5.1.1 Fault Table – ‘faultTable’ 1.3.6.1.4.1.17384.1.1.2.2

oid = "1.3.6.1.4.1.17384.1.1.2.2.1.ifIndex"

portMonitoringFaultStatus: fault status for each port.

Type: integer

1 = fault has occurred

2 = no fault has occurred

Access: read-only

5.1.5.2 RapidRing Group – 1.3.6.1.4.1.17384.1.1.3

oid = "1.3.6.1.4.1.17384.1.1.3.1.0"

ringEnableStatus: the RapidRing is enabled or disabled

Type: integer

1 = RapidRing is enabled

2 = RapidRing is disabled

oid = "1.3.6.1.4.1.17384.1.1.3.2.0"

switchMode: the switch is a Master or a Slave

Type: integer

1 = Master

2 = Slave

oid = "1.3.6.1.4.1.17384.1.1.3.3.0"

ringID: the group number of the RapidRing

Type: integer

1 = group number 1

2 = group number 2

oid = "1.3.6.1.4.1.17384.1.1.3.4.0"

networkStatus: the network topology

Type: integer

1 = Ring Not Available (because RapidRing is not enabled)

2 = Ring Complete

3 = Ring Incomplete

5.1.6 Message Format for SNMP Operations

Five SNMP operations are used in SNMP version 1 : *get*, *get-next*, *set*, *get-response*, *trap*.

The first four commands are used to send and receive information for managed objects and use the same message format. *Trap* uses a different format discussed in Section 5.1.6.2.

5.1.6.1 Format of Command Messages

Each command message contains a header and a protocol data unit (PDU).

5.1.6.1.1 Message Header

The fields in the message header contain :

Version — 0, indicating SNMP version 1.

Community string — the community string, which authorizes NMS access to the switch.

5.1.6.1.2 Message Protocol Data Unit (PDU)

SNMPv1 PDUs contain a specific command (*get*, *set* etc.) and operands that indicate the object instances involved in the transaction.

The fields in the PDU contain :

PDU type — indicates the command type : *get(0xA0)*, *get-next(0xA1)*, *set(0xA3)*, *get-response(0xA2)*

Request ID — a 4-octet integer used to match response to queries

Error Status — a single octet integer containing a value of zero in a request and the following error status in a response

noError (0): no problem

tooBig (1): the response to your request was too big to fit into one response.

noSuchName (2): an agent was asked to get or set an OID that it can't find; i.e., the OID doesn't exist. It can be used for an unsupported object.

badValue (3): a read-write or write-only object was set to an inconsistent value.

readOnly (4): this error is generally not used.

genErr (5): none of the previous errors.

Error Index — a single octet integer that associates an error with a particular object identifier (OID). Only the response operation sets this field. Other operations set this field to zero.

Variable binding — contains a sequence of OIDs and values

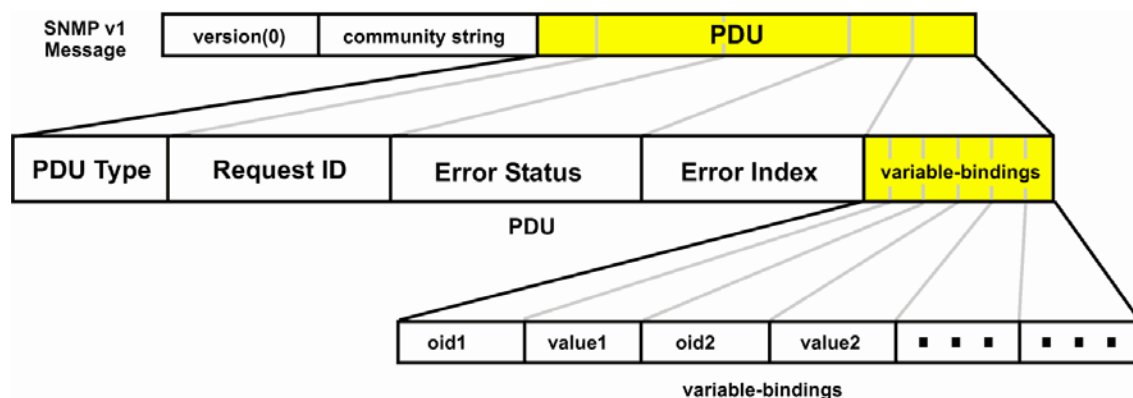


Figure 55 — Format of the Command Message for SNMP Version 1

5.1.6.2 Traps for SNMPv1

5.1.6.2.1 Format of Trap Messages

Each trap message contains a header and a protocol data unit (PDU).

5.1.6.2.2 Trap Header

The fields in the trap message header contain :

Version — 0, indicating SNMP version 1.

Community string — the community string, which authorizes NMS access to the Trap Protocol Data Unit (PDU)

The fields in the PDU contain :

PDU type — 4 (indicates version 1 trap PDU)

Enterprise — Identifies the type of managed object generating the trap. For switch traps, the value is as follows :

Generic Trap — value is *SNMP* (1.3.6.1.2.1.11)

Agent-address — the IP address of the originating agent.

Generic-trap — 0 to 6, indicating the generic trap type. See Section 5.1.6.2.3 for descriptions of the generic-trap types.

Specific-trap — indicates the specific trap type. This field is only interpreted when the generic trap type is 6, *enterpriseSpecific*.

Time-stamp — seconds since last power cycle.

Variable bindings — one or more OIDs (object identifiers) paired with the corresponding values. A variable is an instance of a managed object. These pairings provide more information about the event.

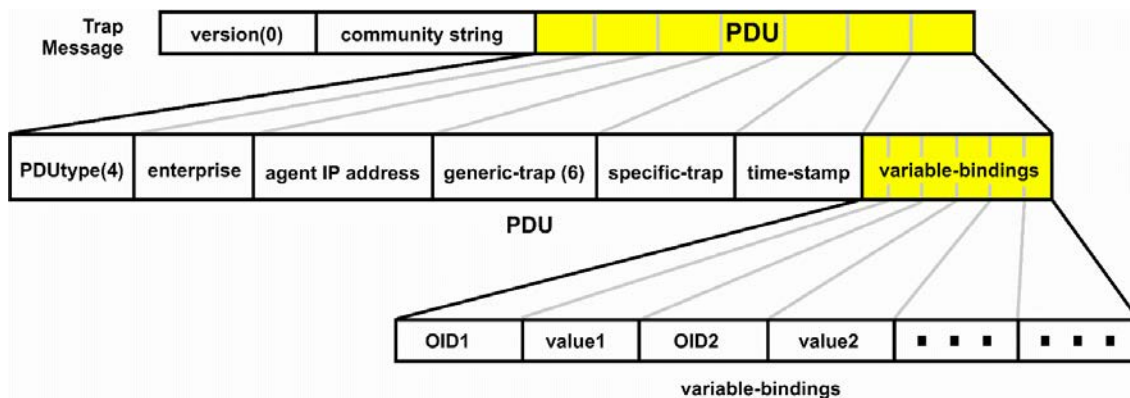


Figure 56 —Format of the Trap Message for SNMP Version 1

5.1.6.2.3 SNMP Generic Traps

enterprise = 1.3.6.1.2.1.11

Generic-trap = 0

coldStart: signifies that the sending protocol switch is reinitializing itself such that the agent's configuration or the protocol switch implementation may be altered.

Generic-trap = 1

warmStart: signifies that the sending protocol switch is reinitializing itself such that neither the agent configuration nor the protocol switch implementation is altered."

Generic-trap = 2

linkDown: signifies that the sending protocol switch recognizes a failure in one of the communication links represented in the agent's configuration.

Generic-trap = 3

linkUp: signifies that the sending protocol switch recognizes that one of the communication links represented in the agent's configuration has come up.

Generic-trap = 4

authenticationFailure: signifies that the sending protocol switch is the addressee of a protocol message that is not properly authenticated.